Generating Under the Influence:
An Adversarial Approach to Modeling Stylistic Influences in Literary Text

Edward Dong

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Science

University of Washington
2024

Reading Committee:
Gina-Anne Levow
Anna Preus

Program Authorized to Offer Degree:
Linguistics

University of Washington

Abstract

Generating Under the Influence:
An Adversarial Approach to Modeling Stylistic Influences in Literary Text

Edward Dong

Chair of the Supervisory Committee:
Gina-Anne Levow
Department of Linguistics

This study proposes, implements, and evaluates a method for simulating *literary stylistic* influence in a text generation system. Specifically, it probes the effects of enhancing next-token prediction training with the introduction of an additional objective, borrowed from the generative adversarial paradigm, for the purpose of administering stylistic pressures. Two crucial adaptations are applied to the typical adversarial objective. First, in order to nudge the model toward preferred styles, the objective is expanded from binary to *multi-label*. And second, in order to cushion the model against the inherent volatility of the adversarial training signal, losses related to non-preferred styles are zeroed out and ignored. All model components are warm-started from the historically pre-trained MacBERTh and fine-tuned on a bespoke corpus of 1950s Anglophone prose fiction. The study additionally devises evaluation metrics grounded in relevant critical and pedagogical literature. The implementation of this socially-adaptive text generation system not only demonstrates a viable approach to modeling peer stylistic influence but may moreover serve as a building block for future research on cultural evolution systems in the literary domain.

TABLE OF CONTENTS

Page

List of Figures iii

List of Tables iv

Acknowledgements v

Chapter 1: Introduction 1
1.1 From Turing to Today 1
1.2 Research Overview 3
1.3 Publication Disclosure 3

Chapter 2: Background 5
2.1 Language Models 5
2.1.2 BERT 5
2.1.3 MacBERTh 6
2.1.4 Can MacBERTh Decode? 6
2.2 The Adversarial Paradigm 7
2.2.1 Gumbel-Softmax 7
2.3 Evaluation of Computational Creativity Systems 8
2.3.1 Style Versus Content 9

Chapter 3: The Corpus 10
3.1 Corpus Curation 10
3.2 Ethical Considerations 12
3.3 Author Overviews 12
3.3.1 Chinua Achebe 13
3.3.2 James Baldwin 15
3.3.3 Shirley Jackson 17
3.3.4 Flannery O'Connor 20
3.3.5 J.D. Salinger 23
3.4 From Corpus to Metrics 25

Chapter 4: The Metrics 26
4.1 Diction 26
4.2 Rhythm 28
4.3 Sentence Complexity 30
4.4 Register 32
4.5 The Strunk and White Metrics 33

4.6 The Strunk and White Metrics, Calculated 34
4.7 Reflection on Automated Metrics 34

Chapter 5: Methodology 36
5.1 The Models 36
5.2 Priming Stage 37
5.3 Adversarial Stage 38
5.3.1 The Intrinsic Style 38
5.3.2 The Influencing Style 40
5.4 Summary of Configurations 42
5.4.1 Model Hyperparameters 42
5.4.2 Training Hyperparameters 43

Chapter 6: Results and Discussion 46
6.1 General Analysis 48
6.2 Error Analysis 50
6.2.1 Type-Token Ratio 50
6.2.2 Average Syllable Count per Single-Breath Segment 51
6.2.3 Percentage of Straight-Through Sentences 51
6.2.4 Exclamation 51
6.3 Summary of Results 51

Chapter 7: Conclusion 53
7.1 Limitations 53
7.2 Future Research 54

Bibliography 55

Appendix A 60

Appendix B 62

Appendix C 67

LIST OF FIGURES

Figure Number Page

4.1 Distribution of Syllable Counts in Sampled Tokens 27
4.2 Distribution of Syllable Counts in Sampled Single-Breath Segments . . . 29

LIST OF TABLES

Table Number		Page
3.1	1950s Anglophone Styles of Fiction	11
4.1	Diction Metrics	27
4.2	Rhythm Metric	29
4.3	Sentence Complexity Metric	31
4.4	Register Metric	32
4.5	Strunk and White Metrics, Formulae	33
4.6	Strunk and White Metrics, Values	34
5.2	Style Discriminator Confusion Matrix	38
5.3.1	Next-Token Prediction Illustration	39
5.4.1	Model Hyperparameters	42
5.4.2A	Priming Stage Configuration	43
5.4.2B	Adversarial Stage Configuration	44
6	Tabulated Results	46
6.2.1	Token Distribution Cosine Similarities Between Corpus Styles	50

ACKNOWLEDGEMENTS

I owe the completion of this voyage to the celestial navigation of my advisor, Gina-Anne Levow; to the resonant foghorns of my reader, Anna Preus; to all those friends and family who held their breath; and to all those whose breath filled my sails.

Chapter 1
INTRODUCTION

"If we are trying to produce an intelligent machine, and are following the human model as closely as we can, we should begin with a machine with very little capacity," Alan Turing proposes in his seminal 1948 paper, "Intelligent Machinery." "Then by applying appropriate interference, mimicking education, we should hope to modify the machine until it could be relied on to produce definite reactions to certain commands. This would be the beginning of the process" (Turing, 1948).

For Turing, "interference" encompasses all the educational, social, and cultural forces that condition an individual's behavior throughout their lifetime, and especially in their formative years. "In fact," he argues, such interference "will be the rule rather than the exception" (Turing, 1948). In other words, if the human model of intellectual development is to serve as any guide, the recipe for machine intelligence calls for continual shaping and reshaping through engagement with other agents.

Seventy-five years after the publication of "Intelligent Machinery," enthusiasm for intelligent machines has precipitously infected the collective human consciousness. Large language models and, most recently, large multimodal models like GPT-4 (OpenAI et al., 2023) and Gemini (Gemini Team et al., 2023) have demonstrated the ability not only to summarize data, ace exams, explain jokes, etc., but also to create convincing cultural artifacts such as stories and poems. It is precisely in relation to this final ability – of not just creation but indeed creativity – that Turing's intuition about interference seems to ring truest.

1.1 From Turing to Today

This insight – that developing the capacity to innovate requires a social environment – has resurfaced in various forms across the field of computational creativity. Notably, in *The Creative Mind: Myths and Mechanisms*, Margaret Boden advances Turing's argument from an epistemological perspective (Boden, 2004). "If, by some miracle, a composer had written atonal music in the sixteenth century, it would not have been recognized as creative. To be appreciated as creative," Boden reasons, "a work of art […] has to be understood in a specific relation to what preceded it." Whereas Turing focuses on casting creativity as a *product* of a provided social

context, Boden seems more focused on interpreting creativity as a qualified *deviation* from that context. Taken together, Turing and Boden's statements establish social context as a necessary prerequisite for creativity: per Turing, necessary for creativity to develop; per Boden, necessary for creativity to count.

Interest in modeling social context is growing. At the 2023 International Conference on Computational Creativity, research teams evaluated algorithms for implementing cultural evolution among a population of image-generating models (Colombo et al., 2023); underscored the role of "creative communities" in shaping creative production (Pease et al., 2023); and posited treating creativity as a multi-agent optimization problem that incorporates influence – both positive and negative – between agents (Luke, 2023). The practical implementation and evaluation of social influence, however, remain challenging.[1]

The present study limits its purview of social influence to the domain of literary style (i.e. to *literary stylistic influence*), further simplifies its conception of literary stylistic influence to the mechanism of emulation, and strives to take on both the implementation and the evaluation of that mechanism. On the implementation side, this study investigates a promising approach to modeling literary stylistic influence. Envision for a moment a human writer's "inner critic": After this writer has read the latest works of their peers and decided what they like *and dislike*, their "inner critic" steps in to monitor the writer's own work – to make sure it sounds more like what they like *and less like* what they don't.

In order to mimic the affinity that a human author may feel for the diverse literary styles of contemporaneous writers, the study introduces a *style discriminator* model whose role is to train a *text generator* model by – in plain second-grade English – "tasting" sentences and responding "yum" or "yuck." Rather than "tasting" a sentence, however, the style discriminator attempts to assign the sentence to one of the available styles; and rather than "yum" and "yuck," of course, the style discriminator responds with numerical values (akin to a probability distribution) indicating how likely the sentence is to belong to each available style. Taking those numerical values as feedback, the text generator can learn to produce sentences that sound more like the agreeable styles and less like the disagreeable ones. Thus, the text generator and style

[1] This is especially true in discrete domains such as text, in which a generative system's outputs change abruptly rather than gradually.

discriminator form a system with the capacity to evolve and hybridize in response to peer stylistic influence.

On the evaluation end, this study devises an evaluative framework grounded in a coherent set of style elements. The framework is designed to isolate *style* as much as possible from *content* (also referred to as *topic*) and used to reveal where the implementation succeeds and where it fails.

1.2 Research Overview

This study is founded on a two-fold objective. First, it endeavors to address a primary research question: *To what extent can the effects of literary stylistic influence be realized by leveraging an adversarial training signal?* And second, it poses the related question: *Which kinds of style features are most learnable under the proposed paradigm, and which are least?*

En route to resolving these questions, this study provides crucial background in Chapter 2 on the relevant types of training signals as well as on the model architectures that will either generate or ingest those signals. Chapter 3 describes and justifies the curation of a literary corpus drawing from the works of five authors from 1950 to 1959. The chapter also gives, for each of the five authors, an overview of the body of criticism he or she inspired; from these overviews, a handful of style elements are extracted and then synthesized into quantifiable style features – the subject of Chapter 4. At last, after the training corpus and the evaluative metrics have been defined, Chapter 5 ventures into the experimental methodology, followed by Chapter 6's analysis of the experimental results.

Ultimately, this study reaches the conclusion that the proposed methodology is capable of engendering influence effects along some but not all of the examined style dimensions. The methodology appears to be most successful at inducing shifts in style dimensions closely related to word choice, such as diction and register. In contrast, the system fails to exhibit influence effects along more abstract dimensions like rhythm and sentence complexity. Depending on their design, the style features belonging to those broader style dimensions may also find their learnability hamstrung by their own sparsity in the training corpus.

1.3 Publication Disclosure

Parts of this thesis are expected to appear, in an abridged format, as a long paper in a special issue of *MDPI Logics* titled *ESSLLI Student Session '23-'24*. Partial content overlap occurs in Chapter 1 (except for Sections 1.2 and 1.3), Chapter 2 (except for Section 2.2.1), Chapter 3 (except for Section 3.3), Section 4.5 from Chapter 4, Chapter 5, Section 6.1 from Chapter 6, and Chapter 7.

Chapter 2
BACKGROUND

This chapter is composed of three sections. The first two – "Language Models" and "The Adversarial Paradigm" – are broadly concerned with the construction of generative models, whereas the last – "Evaluation of Computational Creativity Systems" – discusses their evaluation, particularly in the computational creativity context.

2.1 Language Models

In essence, a language model assigns probabilities to sequences of tokens in a language. These probability distributions provide the basis for text generation. Simply a sequence of tokens, text can be generated via an incremental process of probabilistic sampling from a vocabulary of allowed tokens. This procedure is often referred to as *next-token prediction*.

In contrast to their forebears (such as n-gram language models), neural language models represent token sequences using high-dimensional numerical vectors; it is through operations on these vectors that probability distributions are computed. One of the advantages of the latest major innovation in neural language models, the *transformer*, is its ability to encode within those vector representations of token sequences the linguistic dependencies that exist between any pair of tokens, regardless of the tokens' proximity in the sequence (Vaswani et al., 2017).

2.1.2 BERT

Few of the "early" models built upon the transformer's groundbreaking architectural foundation have shared the enduring impact of BERT, or Bidirectional Encoder Representations from Transformers (Devlin et al., 2018). In a departure from previous transformer models, BERT discards the decoder half of the transformer and leverages only the encoder half; it prioritizes understanding and *representing* text rather than generating new text.

Of particular relevance to this study is the *bert-base-uncased* model (Devlin et al., 2018). This specific BERT variant consists of 12 attention layers, an internal hidden size of 768, and 12 attention heads, for a total of 110 million parameters. While bert-base-uncased itself is a specific pre-trained language model, its architectural hyperparameters have been used as a template for otherwise unrelated models.

2.1.3 MacBERTh

One model that borrows the bert-base-uncased architecture is MacBERTh (Manjavacas et al., 2022), a historically pre-trained encoder model developed on English-language corpora spanning 1450-1950. In contrast to similarly motivated large language models such as the historically *fine-tuned* BERT and Flair models from Hosseini et al. (2021), MacBERTh is unique in having been trained entirely *from scratch* on a temporally delimited corpus.

This study constructs several system components on top of the historically pre-trained MacBERTh model. The selection of MacBERTh is intentional: It permits the researcher to maintain *experimental control* over the system's data sources. Recent studies have illuminated the breadth of copyrighted literary material such as novels (Chang et al., 2023) and poems (D'Souza et al., 2023) that are directly retrievable from large language models like ChatGPT; the implicit knowledge of such material could manifest itself variably in a creative system's behavior, in certain cases compromising the validity of experimental results from downstream tasks.

Thus, while MacBERTh may not yield the same multifaceted linguistic capabilities as a model trained on unbounded English-language data, it offers an important advantage. The researcher is assured that any model built on top of MacBERTh has not seen, and therefore cannot draw upon, examples of post-1950 usage of the English language *outside of any works included in the researcher's own fine-tuning corpus*. The models presented in this study can therefore be thought of as being situated at a precise textual horizon. Crucially, this stringent standard of experimental control lends legitimacy to any claims of what Boden termed "P-creativity" (i.e. production that is *new to the system*)[2] that may arise from future qualitative analysis related to this system (Boden, 2009).

2.1.4 Can MacBERTh Decode?

Models of the BERT family have been fine-tuned not only for sequence-to-value tasks such as sentiment analysis and text classification, but also for sequence-to-sequence tasks like named entity recognition and question answering. The viability of repurposing transformer encoder checkpoints like BERT and RoBERTa (Liu et al., 2019) as *decoders* for sequence generation

[2] as opposed to "H-creativity" – new in the sense of "historical first"

tasks has been studied and verified by Rothe et al. (2020). On a range of benchmarks, decoder components warm-started from a pre-trained BERT checkpoint outperformed both randomly initialized decoders and GPT-2 (Radford et al., 2019). In this study, MacBERTh will be fine-tuned in both text classification and text generation capacities.

2.2 The Adversarial Paradigm

A conventional GAN, or Generative Adversarial Network (Goodfellow et al., 2014), requires two opposed models and a dataset of real artifacts. One of the models, the *generator*, has for its objective the generation of fake artifacts that the other model, the *discriminator*, will fail to distinguish from the real artifacts. The discriminator's objective is not to be fooled: It endeavors to distinguish accurately between real artifacts and the fake ones coming out of the generator. As the generator and discriminator grow concomitantly stronger by training against each other, the generated artifacts are expected to approximate better and better some portion, if not all, of the real distribution.

As will become clear in Chapter 5, this study adapts a conventional GAN for a wholly different purpose – to apply stylistic pressures to a text generator. Specifically, the study leverages a style discriminator to critique the text generator's output. The text generator's objective is to generate artifacts that the style discriminator will judge as likely to belong to predetermined target styles; the style discriminator is a primed multi-class classifier that distinguishes between multiple literary styles (rather than between real and fake).

2.2.1 Gumbel-Softmax

Applying the adversarial paradigm to text generation is not straightforward, however. In contrast to images, where the output space is continuous and differentiable, text involves discrete tokens, which are non-continuous and therefore non-differentiable. More concretely, the generator's output must first be sampled into discrete tokens before it can be evaluated by the discriminator, but sampling is a non-differentiable operation. This non-differentiability interrupts the flow of gradients from the discriminator (which judges the text) to the generator (which has produced the text). With the gradient cables effectively snipped, feedback from the discriminator is prevented from reaching and training the generator.

To circumvent this issue, the Gumbel-Softmax trick (Jang et al., 2016) provides a method for *approximating discrete sampling with a differentiable function*. The Gumbel-Softmax works by introducing noise, sampled from the Gumbel distribution, to the value that the generator outputs for each token in the vocabulary. The randomness of the Gumbel noise effectively stands in for the randomness of the token sampling operation, thereby eliminating entirely the need for that operation. The overall randomness of token sampling has not simply vanished; rather, it has been displaced and relegated to the Gumbel noise.

After the addition of Gumbel noise, the generator's output is converted into a continuous approximation of a categorical distribution using a softmax operation calibrated by a temperature parameter. When the temperature is high, the output is smooth and closer to a uniform distribution; as the temperature decreases, the output becomes sharper and more closely approximates a one-hot vector.

By integrating the Gumbel-Softmax trick, this study overcomes the primary obstacle of gradient blockage in adversarial text generation systems, and thus gradient-based optimization can finally be applied. Chapter 5 situates this key trick in relation to the broader objective of fomenting stylistic shifts in the generator.

2.3 Evaluation of Computational Creativity Systems

Evaluation of the output of creative systems is often imprecise, prejudiced, or missing entirely (Ritchie, 2001) (Jordanous, 2012). Lamb et al. (2018) provide a blueprint for tackling the intractable task of creativity evaluation by breaking it down into mutually-independent aspects, following the 4 P's taxonomy (Person, Product, Process, and Press) introduced by Rhodes (1961). In particular, Lamb et al. define the *Process perspective* as encompassing all theories of "*how* creative products are made" – such as iterative paradigms characterized by a "loop between generation and evaluation." For research motivated by this perspective, including the present study, Lamb et al. suggest constructing an argument that the system is creative because of "how it works" and "what it does and does not do for itself," rather than because of factors like aesthetic value (the Product perspective), audience reception (the Press perspective), etc.

Therefore, this study's approach to evaluation is to demonstrate that the autonomous mechanism does in fact work *as intended* by displaying evidence of the hypothesized style shifts.

2.3.1 Style Versus Content

Measurement of style, however, is far from straight-forward. This study defines style as it is generally understood among linguists (excluding sociolinguists, who attach a different significance to the term): Style refers to *how* things are said, which can vary even when the *what* of the said things remains roughly the same (Pennebaker et al., 2000).

Owing to inevitable correlations between style and content (Wegmann et al., 2022), it may not be possible to fully isolate the former. Nonetheless, previous research has attempted to disentangle stylistic from semantic variation as much as possible. A standard approach is to exclusively examine *stop words* (also known as *function words*), since they encode little to no semantic information. Burrows's Delta (Burrows, 2002) is a classic example of this approach. Commonly applied in authorship attribution tasks, Burrows's Delta leverages stop word frequencies to compute relative stylistic distances between texts.

Similarly, in "Echoes of Power: Language Effects and Power Differences in Social Interaction," Danescu-Niculescu-Mizil et al. (2012) focus on "function word classes, rather than domain-specific substantive content" and furthermore assess the "domain-independence" and generalizability of their approach; they look specifically at eight LIWC (Pennebaker et al., 2001) word classes: "articles, auxiliary verbs, conjunctions, high-frequency adverbs, impersonal pronouns, personal pronouns, prepositions, and quantifiers."

Following the same paradigm, this study hones in on style features that are indifferent to content. In contrast to Danescu-Niculescu-Mizil et al., however, it moves away from LIWC word classes and, instead, identifies a set of 10 style features with high relevance to the study's corpus and to the literary domain, *based on a review of critical and pedagogical literature*. Precisely because they are grounded in the literature, the 10 features developed in this study result in style fingerprints that are intuitive and interpretable.

It is worth mentioning that wholly unrelated paradigms of stylometry exist as well. One particularly interesting approach can be found in the 2020 study, "Interacting with Literary Style Through Computational Tools," in which Sterman et al. reject the use of "explicit features" in favor of a "tacit model of style" based on over 150,000 crowdsourced style judgments (2020).

Chapter 3
THE CORPUS

The corpus and style metrics constitute the first major contribution of this study. This section describes the "1950s Anglophone Styles of Fiction" corpus as well as the curation thereof.

The selection of the 1950s was no accident. The warm-started MacBERTh model that this study takes as its point of departure was trained from scratch on corpora spanning the years 1450-1950: Early English Books Online corpus, Corpus of Late Modern English Texts, Evans Early American Imprints Collection, Eighteenth Century Collections Online, Corpus of Historical American English, and Hansard Corpus (Manjavacas et al., 2022). (These diverse corpora constitute an *inherited dataset* for the models developed in this study.) Thus, the 1950s are precisely the first decade posterior to the historical purview of MacBERTh. The works making up the "1950s Anglophone Styles of Fiction" corpus are therefore guaranteed to be previously unseen by and entirely new to any model warm-started from MacBERTh.

3.1 Corpus Curation

Model training in this study is conducted on a "1950s Anglophone Styles of Fiction"[3] corpus curated by the author. The corpus contains a total of 23,855 sentences spanning five literary styles. Each style is associated principally with an acclaimed novel; most novels are supplemented with one or more short stories penned by the same author and exhibiting similar tone and diction.

The corpus was curated in accordance with the following desiderata:

1. Genre comparability: Works fall in the category of English narrative fiction.[4]
2. Stylistic diversity: Selected novels are representative of distinct literary styles.

[3] In adherence to copyright law, none of the listed works may be publicly released. To request access to the corpus, please contact the author directly.

[4] The only exception is Chinua Achebe's "Polar Undergraduate," a satirical piece penned for a student publication at the University of Ibadan. It is included in this corpus for the purpose of reaching a minimum sentence count for Style A. It may also be worth noting that James Baldwin's novel *Go Tell It on the Mountain* is widely regarded to be semi-autobiographical.

3. Relative contemporaneity: The authors can be assumed to stand upon roughly the same English language substrate. ("Englishes" vary of course – from continent to continent, from community to community, and from individual to individual.)
4. Availability: Works published in the 1950s have not yet been released into the public domain in the United States. However, the majority of the chosen works are available online through educational institutions.

Sentence tokenization was performed carefully by hand. The final sentence counts below reflect the application of two exclusion criteria:

1. Sentences containing any racial slurs
2. Single-word sentences[5] (e.g. "Eleanor!")

Table 1 presents the titles constituting each literary style in the corpus.

Table 3.1: 1950s Anglophone Styles of Fiction

Style	Author	Year	Title	Sentence Count
A	**Chinua Achebe**	1950	"Polar Undergraduate" (Achebe, 1950)	53
		1951	"In a Village Church" (Achebe, 1951)	51
		1952	"The Old Order in Conflict with the New" (Achebe, 1952)	159
		1953	"Dead Men's Path" (Achebe, 1953)	83
		1958	*Things Fall Apart* (Achebe, 1958)	3667
		1959	"The Sacrificial Egg" (Achebe, 1959)	83
			ALL A TITLES	**4096**
B	**James Baldwin**	1953	*Go Tell It on the Mountain* (Baldwin, 1953)	4310
		1957	"Sonny's Blues" (Baldwin, 1957)	991
			ALL B TITLES	**5301**
C	**Shirley Jackson**	1959	*The Haunting of Hill House* (Jackson,	4169

[5] An exception was made for sentences that comprised a single interjection (such as "yes," "no," or "oh") and that were followed by a longer sentence clarifying the cause of the interjection. In such cases, the shorter sentence was prepended to the longer sentence to form a single sample.

			1959)	
			ALL C TITLES	**4169**
D	**Flannery O'Connor**	1952	*Wiseblood* (O'Connor, 1952)	3412
		1955	"A Good Man is Hard to Find" (O'Connor, 1955)	411
		1956	"Greenleaf" (O'Connor, 1956)	542
			ALL D TITLES	**4365**
E	**J. D. Salinger**	1951	*The Catcher in the Rye* (Salinger, 1951)	6729
			ALL E TITLES	**6729**

3.2 Ethical Considerations

Training models on the creative output of human artists is a fraught topic – socially, economically, and ethically. The author does not pretend to co-opt or to essentialize the voice of any writer included in this corpus. Rather, the literary styles are narrowly construed and entirely bound to the specific works listed in Table 1. Hence, models fine-tuned on these styles do not aspire toward (nor do they succeed in) embodying the complex identity of any human writer.

3.3 Author Overviews

The purpose of these biographical and critical snapshots is not to establish the preeminence of the authors or to sing their praises. Rather, they serve to place the authors in their historical, social, and literary contexts, as well as to highlight particular elements of the authors' craft that have been noted by literary scholars. The elements thus skimmed from the critical literature directly informs the development of the style metrics used in this study. As Chapter 4 will make apparent, nearly all of these elements can be translated into automated proxy metrics.

The snapshots serve another crucial purpose. Scrutiny of the corpus' constituent works allows for underlying sensitivities and tensions to surface – and to be acknowledged. To this day, the authors of this corpus remain fiercely beloved, belittled, and provocative; and their works stunningly uplifting, despective, and divisive. Before training models on such a corpus, it is imperative to understand what is in it.

3.3.1 Chinua Achebe

Chinua Achebe was born in 1930 in Ogidi, an Igbophone town in British-ruled Nigeria (Franklin, 2008). Throughout his childhood, he witnessed the complex interactions between native and colonial institutions, which would form an important subject matter in his literary production throughout the 1950s.

In 1948, Achebe entered University of Ibadan (formerly University College), where he wrote short pieces including “Polar Undergraduate” (1950), “In a Village Church” (1951), “The Old Order in Conflict with the New” (1952, appearing in a later version under the title “Marriage is a Private Affair”), and “Dead Men’s Path” (1953). His first novel, *Things Fall Apart* (1958), pushed Achebe into the international spotlight. The following year, he published the short story “The Sacrificial Egg” (1959).

The 1960s would see not only Achebe start a family but also Nigeria gain its independence from British rule, descend into social and political turmoil, and fracture along ethnic lines in a civil war that would subject the country’s Igbo population to genocide and famine. Internally displaced alongside his family, Achebe would shift during this period to poetry as his form and to the war as his primary subject matter (Franklin, 2008).

Literary critics have frequently applied a post-colonial realist analysis to Achebe’s prose fiction. This comes as little surprise: Achebe himself advocated for such a reading. In his 1965 essay “The Novelist as Teacher,” Achebe explains:

> I would be quite satisfied if my novels (especially the ones I set in the past) did no more than teach my readers that their past — with all its imperfections — was not one long night of savagery from which the first Europeans acting on God's behalf delivered them. Perhaps what I write is applied art as distinct from pure. But who cares? Art is important, but so is education of the kind I have in mind. And I don't see that the two need be mutually exclusive. (Achebe, 1965b)

The very real power of fiction to represent – and thereby change – the world became Achebe’s lifelong credo. Literary critic Ato Quayson observes in his essay “Realism, Criticism, and the Disguises of Both: A Reading of Chinua Achebe's *Things Fall Apart* with an Evaluation of the Criticism Relating to It”:

> The confidence Achebe expresses in the realism of his early novels was shared by his critics and led to several critical formations which sought to elucidate the representationalist aspects of his work. The critical tendency that seemed to take

> the novels most evidently as in a one-to-one relationship to reality was that which sought to recover anthropological data about the Igbos from the novels. (Quayson, 1994)

Quayson takes aim at this critical tendency and cautions against "the unconscious urge to read [the African novel] as recording an 'African' reality that comes without mediation" (1994). He proposes the "gap" between African novels and the cultural matrices from which they borrow as an imperative site of literary criticism.

Stylistically, too, Achebe made known how he wished to be read – or, at least, how he did not wish to be read. In another 1965 essay titled "English and the African Writer," Achebe clarifies that he has no intention to use English as a "native speaker" would. Rather, his mission is to blaze the trail for a "new English":

> I feel that the English language will be able to carry the weight of my African experience. But it will have to be a new English, still in full communion with its ancestral home, but altered to suit its new African surroundings. (Achebe, 1965a)

What does this new English look and sound like, in concrete terms? Let's turn to the opening lines of "The Sacrificial Egg":

> Julius Obi sat gazing at his typewriter. The fat chief clerk, his boss, was snoring at his table. Outside, the gatekeeper in his green uniform was sleeping at his post. No customer had passed through the gate for nearly a week. There was an empty basket on the giant weighing machine. A few palm kernels lay in the dust around the machine. Julius went to the window that overlooked the great market on the bank of the Niger. (Achebe, 1959)

The average word count of the sentences in this excerpt hovers just above 11 words. Among these sentences, there are no coordinating conjunctions, no subordinating conjunctions, and only one relative clause ("that overlooked the great market…"). Simple declarative sentences march in and out, one after another. Moreover, the vocabulary is simple. Consider the complete set of descriptive adjectives: "fat," "green," "empty," "giant," and "great." As literary critic B. Eugene McCarthy notes, scholarship on Achebe has attributed this "simple mode of narration and equally simple prose style" to Achebe's reshaping of English to mimic an Igbo or more broadly African style (1985).

McCarthy argues, however, that this "surface simplicity" is actually deceptive. By analyzing *Things Fall Apart* through the lens of rhythm, he brings to light the complexity

beneath Achebe's "simple prose style." In particular, he highlights Achebe's use of repetition (both lexical and syntactic) and "meter." On the author's use of repetition, McCarthy writes:

> The patterning and repetition in Achebe's novel are characteristics of the self-conscious artistry of oral narrative performance, where plot moves by repetition and predictability. [...] Such rhythmic textures establish the narrative method as imitative of the African oral rather than the English "literary" tradition. (McCarthy, 1985)

And on the author's use of "meter," McCarthy performs scansion on sample sentences from *Things Fall Apart* to support his argument that, while "English is considered to be naturally iambic [...], Achebe's prose would seem largely anapestic" (1985). What McCarthy suggests here is that fundamental rhythmic differences exist between Achebe's prose and other prose in English.

These critical observations on sentence complexity, diction, and rhythm in Achebe's fiction style contribute to motivating the metrics described in sections 4.3, 4.1, and 4.2 of Chapter 4, respectively.

3.3.2 James Baldwin

Even as the modern day legacy of his clairvoyant (and maximally tweet-worthy) essays eclipses that of his other works (Walsh, 2018), James Baldwin remains celebrated today as a ground-breaking novelist. Born in Harlem, New York, in 1924, he moved to Paris in 1948 and, shortly thereafter, in 1953, published his first novel – a semi-autobiographical work set in Harlem titled *Go Tell It on the Mountain*.

The rest of the 1950s saw Baldwin establish himself as a talented Black writer and intellectual. Throughout the decade, he would publish numerous novels, essays, and short stories – including "Sonny's Blues" (1957), a fictional story (also set in Harlem) about the place of art and specifically that of music in mid-century African American life. From the 1960s onward, his career would take a sharp turn: Although he continued to write prolifically, Baldwin would also become a cultural celebrity and a polarizing spokesperson for the African American community. In the crucible of the Civil Rights Movement, much of Baldwin's post-1960 production would be pilloried by literary critics who lamented that *activism* had compromised and diluted his *artistic* potential (Francis, 2014).

In her ambitious metacritical project, *The Critical Reception of James Baldwin, 1963-2010*, Conseula Francis describes Baldwin as a writer who resists placement in any one particular tradition (2014). Citing examples of "critical neglect" in major African American anthologies, Francis suggests that Baldwin occupies a "central-yet-still-marginal status in African American letters." At the same time, Francis notes, a paradoxically "wide range of critical schools" have laid claim to Baldwin – "feminist, queer theory, reader response, poststructuralist, ethnic and multicultural, Marxist, psychoanalytic, new critical" – and have at different periods of time dominated the critical discourse surrounding the author.

Francis characterizes *Go Tell It on the Mountain* as "beautiful" and "lyrical" (2014). Stylistically, Baldwin's prose stands out for its musicality – solemn and enrapturing. "Incantatory," critic Hilton Als proclaims in his essay, "The Enemy Within: The making and unmaking of James Baldwin" (1998). Als deftly describes Baldwin's as "a style that blended a full-throated preacherly cadence with the astringent obliquities of a semi-closeted queen." Consider the opening lines of "Sonny's Blues":

> I read about it in the paper, in the subway, on my way to work. I read it, and I couldn't believe it, and I read it again. Then perhaps I just stared at it, at the newsprint spelling out his name, spelling out the story. I stared at it in the swinging lights of the subway car, and in the faces and bodies of the people, and in my own face, trapped in the darkness which roared outside. It was not to be believed and I kept telling myself that, as I walked from the subway station to the high school. And at the same time I couldn't doubt it. I was scared, scared for Sonny. (Baldwin, 1957)

Together with the text's frequent pauses, Baldwin's heavy use of anaphora (i.e., the repetition of "I…") produces the "incantatory" rhythm that Als identifies. That the anaphoric segment comprises the first person pronoun in particular generates a confessional and introspective mood – a "preacherly cadence" indeed. In the following excerpt from *Go Tell It on the Mountain*, the anaphoric segment is "rage and weeping" and its echo "rage [...], weeping [...]":

> It was a sound of rage and weeping which filled the grave, rage and weeping from time set free, but bound now in eternity; rage that had no language, weeping with no voice -- which yet spoke now, to John's startled soul, of boundless melancholy, of the bitterest patience, and the longest night; of the deepest water, the strongest chains, the most cruel lash; of humility most wretched, the dungeon most absolute, of love's bed defiled, and birth dishonoured, and most bloody, unspeakable, sudden death. (Baldwin, 1953)

Baldwin's language here is ornate: "melancholy," "patience," "night," "water," "chains," and "lash" become "boundless melancholy," "bitterest patience," "longest night," "deepest water," "strongest chains," and "most cruel lash." Superlative modifiers crash like burgeoning waves, one after another, carrying the reader up to a terrifying fever pitch of hyperbolic expression. This crescendo is broken finally by two braking commas: "and most bloody, unspeakable, sudden death." This exaggerated grandiosity is perhaps what Als has in mind when, comparing Baldwin to the black writers who would succeed him in the 1970s, he comments that "Baldwin's fastidious thought process and his baroque sentences suddenly seemed hopelessly outdated, at once self-aggrandizing and ingratiating" (1998).

As mentioned earlier, the dimensions of rhythm and sentence complexity are translated into quantitative metrics in Chapter 4, sections 4.2 and 4.3.

3.3.3 Shirley Jackson

Critic Angela Hague calls Shirley Jackson "a quintessential writer of the 1950s whose work dramatizes the concerns and fears of that decade" (2005). Born in 1916 in San Francisco, California, Jackson's six completed novels and over 100 short stories channeled "the loneliness, isolation, and psychological problems of suburban [middle-class] housewives":

> By focusing on her female characters' isolation, loneliness, and fragmenting identities, their simultaneous inability to relate to the world outside themselves or to function autonomously, and their confrontation with an inner emptiness that often results in mental illness, Jackson displays in pathological terms the position of many women in the 1950s. [...] Jackson's characters possess neither gyroscope nor radar; lacking a core of identity forces them to seek meaning and direction in the world outside themselves, but their inability to relate to and communicate with others and their fear of unfamiliar environments create the panic and paranoia that descend upon them when they venture beyond the dubious safety of their domestic environment. (Hague, 2005)

Jackson herself suffered from anxiety, agoraphobia, and a "problematic" relationship with her mother that would remain unresolved up through Jackson's passing in 1965 (Rubenstein, 1996). Rubenstein describes Jackson's novel protagonists as "typically at the pivotal age between childhood and womanhood" and "particularly anxious and ambivalent about their relationships with their mothers":

> In the first two novels, the mother is invasively present in the daughter's life; in each of the four succeeding novels, the mother is dead but no less powerfully

> present. In fact, the mother's absence becomes a haunting presence that bears directly on the daughter's difficult struggle to achieve selfhood as well as to express her unacknowledged rage or her sense of precariousness in the world. (Rubenstein, 1996)

The Haunting of Hill House, published in 1959, was Jackson's fifth novel. Rubenstein agrees with the critical consensus classifying *The Haunting of Hill House* as Gothic, particularly given the protagonist's "anxieties about self-hood and entrapment, represented through bizarre or exaggerated events that may or may not be explained as manifestations of the [...] central character's troubled imagination" (1996).

While Rubenstein makes use of the term "female Gothic," other critics have variously described Jackson's style as quotidian Gothic (Ashton, 2018), new American Gothic (Parks, 1984), or simply Gothic (Parks, 1984). Hilarie Ashton provides this justification for the "quotidian Gothic" label:

> In *The Haunting of Hill House* (1959), her penultimate novel, Jackson illustrates her mastery of what I call the quotidian Gothic, a term I use to hint at the deep interplay of repression, fear, and disgust inside a more or less "normal" world as portrayed in a nonetheless indisputably Gothic text. Her melding of two different kinds of realities, filtered through an entirely unreliable narrator, is what has made the monstrous magic of contemporary authors like Victor LaValle and Joyce Carol Oates (among others) possible, melding the Gothic genre firmly with the postmodern. [...] The horror in the novel lies largely in the space between Eleanor's [the protagonist's] mind and the external world (and the people in it) rather than in the direct appearance of horrific things, and those more classically monstrous things that do appear are pretty clearly arbitrated by that open space. Everything normal is upended in Hill House, both by the narrative voice and by the plot: Jackson leads her reader into a world where accepted, connective reality rapidly unravels as filtered through the lens of an unstable third-person omniscient narrator [...] (Ashton, 2018)

This "unstable third-person omniscient narrator" is precisely what permits Jackson to make many of her signature – and more experimental – stylistic moves. Ashton pinpoints passages in *The Haunting of Hill House* wherein "Jackson offers up Eleanor's perceptions expressly without the mediation of the narrator, using a variety of narrative devices, from verb tenses to punctuation and formatting choices [...]" (2018). Take, for example, this extravagant fantasy that erupts while Eleanor speeds away to Hill House:

> The road, her intimate friend now, turned and dipped, going around turns where surprises waited—once a cow, regarding her over a fence, once an incurious dog—down into hollows where small towns lay, past fields and orchards. On the

> main street of one village she passed a vast house, pillared and walled, with shutters over the windows and a pair of stone lions guarding the steps, and she thought that perhaps she might live there, dusting the lions each morning and patting their heads good night. Time is beginning this morning in June, she assured herself, but it is a time that is strangely new and of itself; in these few seconds I have lived a lifetime in a house with two lions in front. Every morning I swept the porch and dusted the lions, and every evening I patted their heads good night, and once a week I washed their faces and manes and paws with warm water and soda and cleaned between their teeth with a swab. Inside the house the rooms were tall and clear with shining floors and polished windows. A little dainty old lady took care of me, moving starchily with a silver tea service on a tray and bringing me a glass of elderberry wine each evening for my health's sake. I took my dinner alone in the long, quiet dining-room at the gleaming table, and between the tall windows the white panelling of the walls shone in the candlelight; I dined upon a bird, and radishes from the garden, and home-made plum jam. When I slept it was under a canopy of white organdie, and a nightlight guarded me from the hall. People bowed to me on the streets of the town because everyone was very proud of my lions. When I died… (Jackson, 1959)

Time is "strangely new and of itself" not only for Eleanor, but also for the reader: While the frame narrative stays in the simple past ("she passed a vast house"), Eleanor's reverie springs from past conditional ("she might live there") to simple present ("Time is beginning this morning") to present perfect ("I have lived a lifetime") to past imperfect ("Every morning I swept the porch") to simple past ("When I died"). Dissatisfied with the life she has led thus far, Eleanor attempts to reclaim the past by rewriting it, notably with the timelessness of the imperfective aspect; she seems ready to give herself up in favor of this alternative vision of herself.

The passage is experimental not only for its use of verb tense and aspect but also for the narrative slippage that it exhibits. Narrative control is ceded to Eleanor as a consequence of "[t]he lack of italics, quotation marks, or any other assignation to demarcate thoughts," leading to "Eleanor's unmediated subjective entrance onto the scene" (Ashton, 2018). The unmediated appearance of the first person pronoun in Jackson's novel distinguishes her from other authors in the corpus – in particular from Flannery O'Connor who tended to convey characters' interior attitudes in the third person, even in moments of narrative slippage.

(Neither the instability in narrative tense nor that in narrative perspective is easily observable within a single stand-alone sentence, which is the unit of analysis in this study. Such dimensions of style are therefore excluded from the study's analytical lens. Nonetheless,

Jackson's prose does project a striking stylistic fingerprint along the dimensions that *are* included in the analytical lens, as will become apparent in Chapter 4.)

3.3.4 Flannery O'Connor

Flannery O'Connor was born in Savannah, Georgia, in 1925. "Except for a creative writing session at the University of Iowa she has spent most of her thirty-six years in her native Georgia," literary critic Melvin J. Friedman wrote in 1962. Situating her within the "southern gothic" tradition, Friedman describes O'Connor as showing "the same devotion to birthplace as so many other southern writers. The settings of her stories and novels are either Georgia or Tennessee, often backwoods or rural areas" (1962).

Two years later, O'Connor would succumb to lupus erythematosus. She was diagnosed with the auto-immune disease in 1951, and for the last thirteen years of her life, "she wrote from a farm in Milledgeville, Georgia, cared for by her mother, Regina, who tended both the farm and her daughter," wrote the author Joseph Torchia. "She had the strength to write only three hours a day" (Torchia, 1996).

O'Connor published her first novel, *Wise Blood*, in 1952, and a second in 1960; she was working on a third when she passed. She would be celebrated, however, not as a novelist but as a master of the short story, of which she penned over 30. Alluding to both of the novels published during her life, Friedman gives the following characterization:

> [O'Connor's] rhythm of writing shorter fiction and then, on occasions, reshaping it slightly to suit the needs of a novel and, on others, allowing it to remain intact gives one the impression (and this has frequently been said before about her) that she is essentially a story writer who has twice strayed into the novel form. (Friedman, 1962)

O'Connor's most acclaimed short story, "A Good Man is Hard to Find," was first published in a 1953 anthology. "Greenleaf," on the other hand, was first published in 1956 in *The Kenyon Review*.

O'Connor's fiction has provoked – short stories and novels alike – three enduring lines of criticism: concerning the author's stern Catholic theology, her commitment to the grotesque, and her views on race. Critic Robert Ellsberg quotes a letter in which O'Connor identifies herself as a "Christian Realist" (2005). He paints a compelling portrait of an artist who "lived in the presence of certain theological truths […] as solid as the laws of physics":

> In contrast to many other Catholic writers - whether J.F. Powers, Graham Greene, or Mary Gordon - she did not draw on her faith to supply the "décor" of her stories. Her subjects were more often backwoods fundamentalists, or the "good country people" she encountered among her neighbors. It was not the settings of her stories but her overall point of view that defined her as a Catholic artist. Simply, Flannery O'Connor saw the world in light of the great doctrines of the church: the Fall, the Incarnation, Redemption. That was the world she depicted in her fiction. But as her posthumously published letters make clear, that was also the world she inhabited. In the way she wrote as well as the way she lived she has made it possible for many others to experience life from the standpoint of what she called the central Christian mystery: "that it has, for all its horrors, been found worth dying for." (Ellsberg, 2005)

In his influential *Nightmares and Visions: Flannery O'Connor and the Catholic Grotesque*, critic Gilbert H. Muller defines a "Catholic grotesque" which is grotesque in that disturbed characters are caught in a bewildering, senseless, and absurd reality; and Catholic in that the cause of this absurdity – alienation from God – also suggests its cure (even if not actually attainable by O'Connor's protagonists) (1972). "Catholic grotesque" is perhaps another articulation of the "paradoxical" quality that Friedman elucidates but does not name:

> [H]er novels and short stories [...] abound in sordidness and poverty and yet maintain a delicate aesthetic balance on the side of gentility and religious affirmation. (Friedman, 1962)

On O'Connor's racial views, it will suffice here to note that the author has attracted attention for the equivocal attitude she expressed toward African Americans in her personal life and in her writing. Julie Armstrong's essay "Blinded by Whiteness: Revisiting Flannery O'Connor and Race" (2001) offers an informed treatment centered around the question, "Why do O'Connor's stories so often prompt readers to ask about her racial views?" – a compelling question that will not be treated here.

Given O'Connor's "unwillingness to make pronouncements about her craft," little is known about how the author perceived her own style. Friedman characterizes her style broadly as "unpretentious":

> Her novels and stories are in every sense traditionally constructed and make no use of the experimental suggestions of a Joyce, a Proust, a Faulkner, or even a Styron. Her work is usually completely faithful to chronology, with no attempt at reproducing an atmosphere of psychological time. In short, her fiction bears no relation whatever to the so-called art novel. (Friedman, 1962)

Friedman claims, by way of example, that O'Connor "avoids the more experimental possibilities of stream-of-consciousness fiction, as she does everywhere else in her work, but does try to approximate the workings of the characters' minds if only through the indirect third person" (1962). The following passage from *Wise Blood* narrates protagonist Hazel Motes' thought process at the symbolic start of his journey and bears out Friedman's observation:

> The train jolted and stirred him half-awake again and he thought, there must have been twenty-five people in Eastrod then, three Motes. Turning in the road, he saw in the dark the store boarded and the barn leaning and the smaller house half carted away, the porch gone and no floor in the hall. It had not been that way when he was eighteen years old and had left it. Then there had been ten people there and he had not noticed that it had got smaller from his father's time. He had left it when he was eighteen years old because the army had called him. He had thought at first he would shoot his foot and not go. He was going to be a preacher like his grandfather and **a preacher can always do without a foot. A preacher's power is in his neck and tongue and arm.** [...] He knew by the time he was twelve years old that he was going to be a preacher. Later he saw Jesus move from tree to tree in the back of his mind, a wild ragged figure motioning him to turn around and come off into the dark where he was not sure of his footing, where he might be walking on the water and not know it and then suddenly know it and drown. Where he wanted to stay was in Eastrod **with his two eyes open, and his hands always handling the familiar thing, his feet on the known track, and his tongue not too loose.** [emphasis mine] (O'Connor, 1952)

While the narrative does appear to reflect the sudden turns and discontinuities of the protagonist's thought process as he hovers at the edge of sleep, it does not stray far from the omniscient third person perspective. Segments where one might plausibly argue that the narrative drifts away from the omniscient narrator's voice and into the protagonist's less formal idiom are bolded (emphasis mine); these segments are few and exhibit none of the "more experimental possibilities of stream-of-consciousness fiction," such as a tell-tale loss of grammaticality (Friedman, 1962). That said, O'Connor's use of free indirect discourse is much less ambiguous: that "a preacher can also do without a foot" is unlikely to be the omniscient narrator's opinion, and what constitutes the "familiar thing" is not at all clear to the reader, but it probably is to Hazel Motes.

(As in Jackson's case, O'Connor's prose possesses a distinct profile under this study's analytical framework, even though none of Catholic theology, commitment to the grotesque, or views on race is suitable for direct translation into a *style* metric. The three critical foci seem more amenable perhaps to translation into *content* metrics.)

3.3.5 J.D. Salinger

In his 1997 social history of *The Catcher in the Rye*, American studies scholar Stephen J. Whitfield makes the claim that "[n]o American writer over the past half-century has entranced serious young readers more than Salinger" (1997). Born in 1919 in Manhattan, New York, J.D. Salinger had just established himself as a talented short story writer by the time he published *The Catcher in the Rye*, his first novel, in 1951. It would garner enduring popular acclaim and critical attention, although how exactly it did so has been, in Whitfield's words, "a bit mystifying" (1997). Perhaps the reason lies in the touching evocation of a presumably universal teenage disenchantment, as Carol and Richard Ohmann interrogate (and complicate) in their essay "Reviewers, Critics, and *The Catcher in the Rye*":

> For us, as for almost all readers, Holden's [the protagonist's] sensitivity is the heart of the book, that which animates the story and makes it compelling. Events are laden with affect for Holden. He cannot speak of an experience for long in a neutral way, apart from judgment and feeling. (Ohmann et al., 1976)

Or perhaps the reason lies in the evocation of a specific post-war "social reality" and its contradictions, as Ohmann et al. ultimately propose:

> The Catcher in the Rye is among other things a serious critical mimesis of bourgeois life in the Eastern United States, ca. 1950--of snobbery, privilege, class injury, culture as badge of superiority, sexual exploitation, education subordinated to status, warped social feeling, competitiveness, stunted human possibility, the list could go on. [...] Iran and Korea and the hard-pressed New York school system express the hegemony of Holden's class, as do Broadway and Pencey and Stradlater. Salinger's novel makes no reference to the economic and military scope of that class's power, but the manners and institutions he renders so meticulously are those of people who take their power for granted, and expect their young to step into it. (Ohmann et al., 1976)

At the same time, *The Catcher in the Rye* had its fair share of detractors. Critics decried its insufficient realism (Ohmann et al., 1976), "pedestrian content" (Roemer, 1992), unpatriotic sentiment (Rosen, 1977), morally corrupting influence (Whitfield, 1997), and "obscenity" (Costello, 1959). The quality highlighted by the last objection, especially, is more feature than flaw. The novel is narrated in the first person by a freshly expelled 17-year-old, and Salinger skillfully matches "syntax and idiom to that choice" (Ohmann et al., 1976). As author and scholar Donald P. Costello playfully remarks in "The Language of 'The Catcher in the Rye',"

"No one familiar with prep-school speech could seriously contend that Salinger overplayed his hand in this respect" (1959). Costello suggests that critical reception for the novel's diction was generally positive:

> Most critics who looked at *The Catcher in the Rye* at the time of its publication thought that its language was a true and authentic rendering of teenage colloquial speech. Reviewers in the *Chicago Sunday Tribune*, the *London Times Literary Supplement*, the *New Republic*, the *New York Herald Tribune Book Review*, the *New York Times*, the *New Yorker*, and the *Saturday Review of Literature* all specifically mentioned the authenticity of the book's language. (Costello, 1959)

More specifically, Salinger's text was perceived by contemporaries as an "accurate rendering of the informal speech of an intelligent, educated, Northeastern American adolescent." Costello hones in on "the trite repetitive vocabulary which is the typical lot of his [Holden's] age group." For example:

> There are two major speech habits which are Holden's own, which are endlessly repeated throughout the book, and which are, nevertheless, typical enough of teenage speech so that Holden can be both typical and individual in his use of them. It is certainly common for teenagers to end thoughts with a loosely dangling 'and all,' just as it is common for them to add an insistent 'I really did,' 'It really was.' But Holden uses these phrases to such an overpowering degree that they become a clear part of the flavor of the book; [...] (Costello, 1959)

On the former idiosyncrasy, he writes, "Holden's 'and all' and its twins, 'or something,' 'or anything,' serve no real, consistent linguistic function. They simply give a sense of looseness of expression and looseness of thought." Costello proceeds to identify other idiolectic expressions favored by the narrator, including *for God's sake*, *goddam*, *for Chrissake*, *ass*, *as hell*, *crap*, *crazy*, *killed*, *lousy*, *pretty*, *crumby*, *terrific*, *quite*, *old*, and *stupid* (1959). These expressions, many of them interjections or qualifiers, stand out for their stylistic rather than semantic contribution to the text.

Lexicon operates hand in hand with prosody to create the narrator's style. Costello does not neglect the text's prosodic elements:

> The structure of Holden's sentences indicates that Salinger thinks of the book more in terms of spoken speech than written speech. Holden's faulty structure is quite common and typical in vocal expression. (Costello, 1959)

Costello's emphasis on the vocal quality of the novel is significant, not least because it hints at an analytical commonality with Achebe and Baldwin's work, even though each of the three authors had his own literary background and motives.

The measurement of register, diction, and rhythm – stylistic elements underscored by literary criticism of *The Catcher in the Rye* – is the subject of sections 4.4, 4.1, and 4.2, respectively, of the next chapter.

3.4 From Corpus to Metrics

The review of critical literature pertaining to the corpus works surfaces a wealth of candidate stylistic elements for consideration. Some of those stylistic elements come into existence only at a higher *discourse level* of analysis; these are necessarily ignored, as the output size of the text generation system explored in this study is limited to stand-alone sentences, rather than sentence pairs, paragraphs, or even longer units. Other stylistic elements can be fully observed at the sentence level. Chapter 4 will detail a set of metrics related to these elements.

Additionally, five metrics will be introduced that are not directly motivated by the critical literature. Nonetheless, the reader may detect certain resonances between the summaries of this chapter and the metrics of this latter set, such as the adjective and superlative metrics.

Chapter 4
THE METRICS

This study explores a total of 10 metrics, targeting dimensions of literary style from rhythm to part-of-speech preferences. By design, the metrics look at topic-invariable qualities, rather than fixating on individual open class words. They make use of pre-existing resources – Python libraries such as re, Syllables, better-profanity, and Natural Language Toolkit[6] (Bird et al., 2009) – for automatic computation, rather than requiring human annotation. They operate – given constraints on model output in the second part of this study – on the scale of independent sentences, rather than across longer discourse contexts. And finally, they are motivated by critical and pedagogical literature related to the corpus.

In each subsection below, one or more metrics are described, applied to the five corpus styles, and connected when possible to claims made in the critical literature. This evaluation of the corpus styles sets up the necessary foundation for judging the success of the model, as well as for delivering a verdict on the research hypothesis.

4.1 Diction

Diction refers to a narrator's word choice. What kinds of words does the author select in order to create tone (irreverent, bewildered, suspenseful, hallowed, formal, etc.) or produce other effects on the reader? Two metrics are computed with the intent of capturing some facet of diction: type-token ratio and percentage of "long" tokens (defined as spanning 3 or more syllables).

Sentences from each style in the corpus were tokenized using *white-space tokenization* to produce a candidate list of non-unique tokens belonging to the style. Fictional character names (such as *Okonkwo*, *Eleanor*, and *Greenleaf*) were excluded based on the reasoning that they 1) appear with unduly high frequency and 2) reflect what is real within a fictional world rather than a narrator's word choice. Nearly 400 tokens were excluded on this basis.

24,000 tokens were then selected at random without replacement from each style. Type-token ratio is computed by counting the number of unique lower-cased tokens and dividing that count by the total number of sampled tokens, i.e. 24,000. The percentage of tokens with 3 or

[6] referred to henceforth as *NLTK*

more syllables was computed using the Syllables package over the same 24,000 tokens for each style.

Table 4.1: Diction Metrics

Metric	**A**	**B**	**C**	**D**	**E**
Type-Token Ratio	0.137	0.127	0.143	0.128	0.102
% of 3+ Syllable Tokens	8.925	6.613	9.05	6.15	5.992

The two metrics in Table 4.1 above demonstrate substantial agreement. Style C has the highest Type-Token Ratio as well as the highest Percentage of 3+ Syllable Tokens, followed by Style A. Style E has the lowest value for both metrics. Styles B and D fall in the middle with nearly equal Type-Token Ratios.

This data bears out several key claims from the critical literature. Styles C and E each draw from a single novel. In contrast to the “unstable” (Ashton, 2018) third-person omniscient narration of *The Haunting of Hill House* that slips in and out of the protagonist’s “troubled” (Rubenstein, 1996) perspective, the narration of *The Catcher in the Rye* is remarkably stable (even if the narrator himself is not). Recall Costello’s characterization of Holden’s adolescent vocabulary as typically “trite” and “repetitive” (1959).

To provide the reader with another view into the Percentage of 3+ Syllable Tokens metric, Figure 4.1 plots the distribution of token lengths, using the same 24,000 tokens per style previously sampled.

Figure 4.1: Distribution of Syllable Counts in Sampled Tokens

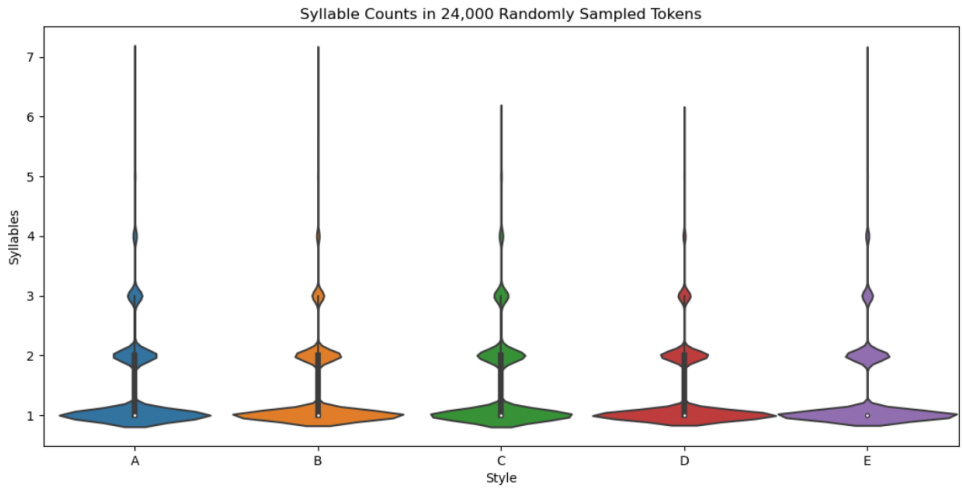

4.2 Rhythm

Rhythm refers to a text's underlying patterns of stresses and pauses. Recall that McCarthy describes Achebe's prose as "largely anapestic" – in defiance of the English language's "natural" tendency toward iambic rhythms (1985). McCarthy's observation hints at fascinating possibilities for the application of metrical analysis to prose in the service of quantifying its rhythmic qualities; however, the pursuit of these possibilities will be left to future research. It may well require the development of more efficient scansion algorithms to account for longer inputs (such as sentences or even multi-sentence paragraphs) and the much larger search spaces that these would entail.

Rather than stress patterns, the following metric aims to capture some tendency in a text's pause patterns. How many syllables, on average, does the reader amble through before being forced to pause? Or, as an alternative way to frame this question, how often is the reader pausing? Although the rhythmic interpretation of any text may differ widely from reader to reader, punctuation is taken here to be a practical indicator of how its author intended it to sound.

To compute this metric, all sentences in the corpus are first split on the positions of commas, semicolons, colons (excluding the time-telling case), ellipses, exclamation marks, question marks, parentheses, and em dashes into "single-breath" segments. (Naturally, instances of these punctuation marks appearing alone or in combination at the beginning or end of a sentence are ignored.) Next, 6000 such single-breath segments are randomly sampled without replacement from each style. Finally, syllable counts for the sampled segments are estimated using the Syllables package and aggregated using a simple average.

Table 4.2: Rhythm Metric

Metric	**A**	**B**	**C**	**D**	**E**
Average Syllable Count per Single-Breath Segment	12.614	8.043	8.699	11.687	8.405

As shown in Table 4.2, styles A and D have a substantially higher average syllable count than the other three styles (roughly 4 syllables longer and 3 syllables longer than the next highest average, respectively). Style B exhibits the shortest average syllable count.

Style B's low value by this metric is, of course, to be expected. It provides quantitative support for the casual observation made earlier that, in at least two excerpts drawn from Style B, commas appear frequently and regularly to produce what Als called a "preacherly cadence" (1998). Similarly, Style A's high value by this metric provides quantitative support for the casual observation that, in the analyzed excerpt from Style A, simple declarative sentences followed one another in what McCarthy deemed a "simple prose style" (1985).

Once again for the reader's benefit, Figure 4.2 offers a visualization of the distribution of syllable counts in the 6000 segments sampled per style.

Figure 4.2: Distribution of Syllable Counts in Sampled Single-Breath Segments

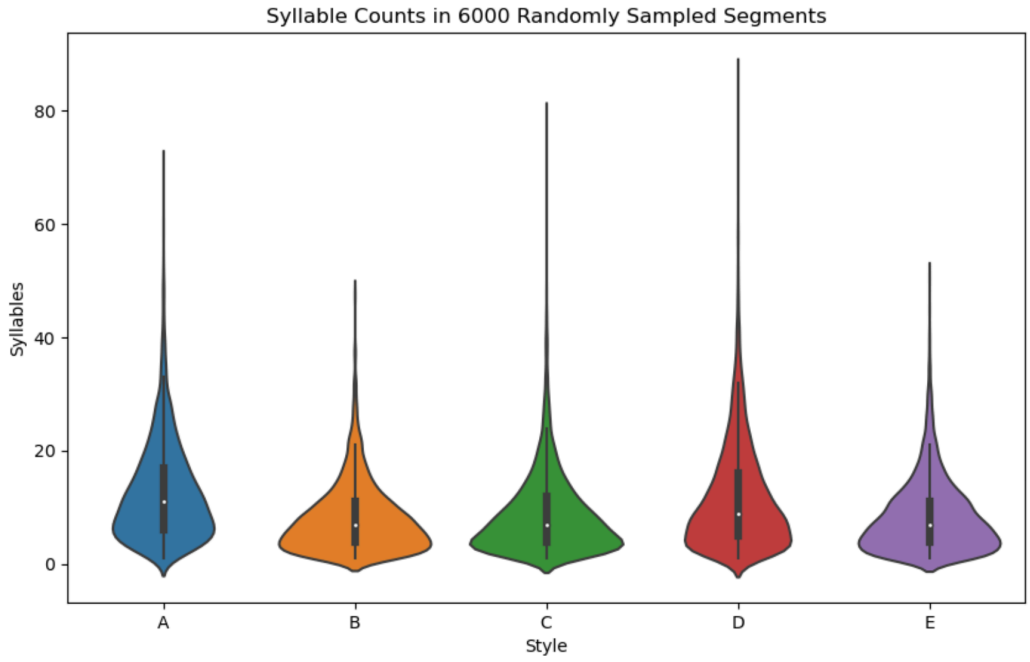

4.3 Sentence Complexity

This metric aims to estimate the frequency of "straight-through" sentences in a literary style. To qualify as "straight-through" in the context of this study, a sentence should be declarative (rather than interrogative, exclamatory, or imperative), simple (rather than compound or complex), and pauseless. In order to detect such sentences and to distinguish them from other sentence types, this metric leverages three rough heuristics:

1. The final stop to the sentence is indicated by a period. This can be true even if the sentence is wrapped in quotation marks, parentheses, etc.
2. The sentence contains at least one noun and exactly one verb conjugated in the present or past tense.
3. The sentence contains no pausing punctuation marks in an internal position. A punctuation mark is considered to appear in an internal position if there exist one or more preceding tokens and one or more succeeding tokens. The set of pausing punctuation marks is the same as that used by the rhythm metric: commas, semicolons, colons (excluding the time-telling case), ellipses, exclamation marks, question marks, parentheses, and em dashes.

The third heuristic operates on the assumption that internal pauses – even when they do not serve their common function of demarcating a coordinate or subordinate clause – still do

alter a reader's perception of the complexity of a sentence. Consider the opening line of Baldwin's short story "Sonny's Blues": "I read about it in the paper, in the subway, on my way to work" (1957). Although technically a simple declarative sentence satisfying heuristics 1 and 2, this line does not feel simple. Each of the prepositional phrases "in the paper," "in the subway," and "on my way to work" is lent gravity and emphasis by Baldwin's comma-work; the distinction between the three separate pieces of information conveyed by the phrases seems to sharpen and solidify. This sentence pretends to present multiple pieces of information, not just one. It is masquerading as a compound sentence, and the reader's eyes and ears are fooled. Heuristic 3 captures precisely this effect.

In the computation of this metric, 3000 sentences are randomly sampled without replacement from each style. Next, the three heuristics are applied to each sentence. (The second heuristic is implemented using NLTK's part-of-speech tagger.) The final score for each style is the number of sentences that satisfy the heuristics over the total number of sentences (i.e. 3000), expressed as a percentage. Final scores are displayed in Table 4.3.

Table 4.3: Sentence Complexity Metric

Metric	**A**	**B**	**C**	**D**	**E**
% of Straight-Through Sentences	27.3	13.933	12.033	19.567	26.067

Only within styles A and E are more than a quarter of sentences estimated to be straight-through. The lowest percentages belong to styles C and B, which exhibit roughly half the straight-through frequency of styles A and E.

That Style A holds the top spot here seems to echo again the critical appraisal of Achebe's prose. Together, the three heuristics give shape to the author's "simple mode of narration" in a computational sense (McCarthy, 1985). On the other end, the "unstable third-person" mode of narration – arguably the most experimental of the corpus – identified by Ashton (2018) in her analysis of *The Haunting of Hill House* lands Style C in the bottom spot. Last and second least, Style B's similarly low straight-through percentage should hardly come as

a surprise to critics like Als who noted Baldwin's penchant for "baroque," amply adorned sentences (1998).

4.4 Register

In the study of literature, register refers to how formal or polite (toward the reader) a text sounds. Like the other elements considered thus far, register will be assessed somewhat reductively, on the basis of a particular facet, rather than holistically. The metric described in this section looks specifically for the presence of *profanity*, which is assumed to decrease a text's level of formality and politeness.

The computation of this metric reuses the same 3000 sentences sampled from each style during the evaluation of the previous metric (Percentage of Straight-Through Sentences). It leverages better-profanity, a wordlist-based profanity detection package, to judge whether each sample contains profanity or not. A percentage of profanity-containing sentences, shown in Table 4.4, is then calculated for each style.

Table 4.4: Register Metric

Metric	**A**	**B**	**C**	**D**	**E**
% of Sentences Containing Profanity	4.4	7.467	1.033	3.3	12.3

Roughly 1 in 8 sentences from Style E is estimated to contain profanity, whereas the next highest frequency, belonging to Style B, is roughly 1 in 13. The lowest frequency belongs to Style C, with profanity present in just over 1 in 100 sentences.

Although this metric ignores some of the specific idiolectic patterns (such as "and all") of the teenage first-person narrator of Style E, it nonetheless succeeds in separating Style E from the rest of the pack. Concretely, the metric characterizes Style E as the most profane; and as a correlate, one may make a softer inference that it is also the most informal. This result echoes observations from the critical literature – after all, Style E is the only style that critics describe as "informal" (Costello, 1959).

4.5 The Strunk and White Metrics

The remaining five metrics are inspired by an influential guide to composition in the English language – titled *The Elements of Style* – that was published in its complete form in 1959 (Strunk, Jr., et al.). Although *The Elements of Style* is a prescriptive rather than a descriptive text, and although it concerns itself with college-level expository prose rather than creative fiction, it nonetheless serves as a useful indicator of the stylistic elements of prose writing that were being talked about around the mid-century. These elements were "in the water," so to speak, and likely understood by professional writers and critics.

Thus, this study draws upon *The Elements of Style* as one approach for isolating specific stylistic features that would have been salient during the 1950s. Furthermore, while the principles themselves are prescriptive (or proscriptive, as is often the case), the metrics constructed from those principles are purely *descriptive*. The features mentioned by the principles, such as adjective or qualifier usage, serve in this study as instruments for describing sentences, rather than for critiquing them. After all, as Strunk himself concedes while introducing *The Elements of Style*, "the best writers sometimes disregard the rules of rhetoric" (Strunk, 1920).

Table 4.5 lists the five metrics motivated by *The Elements of Style*, along with the particular Strunk and White "principles" from which they derive.

Table 4.5: Strunk and White Metrics, Formulae

Metric	**Motivating Principle**	**Formula**
Negation	"Put statements in positive form"	Avg. count of negation tokens per 100 sentences
Adjective	"Write with nouns and verbs"	Avg. ratio (by sentence) of adjectives to nouns
Adverb	"Write with nouns and verbs"	Avg. ratio (by sentence) of adverbs to verbs
Exclamation	"Do not attempt to emphasize simple statements by using a mark of exclamation"	Avg. count of exclamation marks per 100 sentences
Superlative	"Do not overstate"	Avg. count of superlatives per 100 sentences

4.6 The Strunk and White Metrics, Calculated

Table 4.6 displays the Strunk and White metrics computed over the same sampling of 3000 sentences per style that the Percentage of Straight-Through Sentences and Percentage of Sentences Containing Profanity metrics use.

Table 4.6: Strunk and White Metrics, Values

Metric	**A**	**B**	**C**	**D**	**E**
Negation	*17.067*	**41.6**	23.867	31.267	27.133
Adjective	**0.222**	0.161	0.186	*0.155*	0.189
Adverb	0.336	0.384	0.42	*0.307*	**0.468**
Exclamation	1.1	3.3	*0.333*	**4.667**	1.633
Superlative	**2.933**	2.0	2.233	*1.233*	2.167

These metrics are computed using NLTK's part-of-speech tagger along with a custom token set for the Negation metric (written here as regular expressions, where "\b" indicates a word boundary): "\bno\b", "\bnot\b", "\bcannot\b", "\bnone\b", "\bnothing\b", "\bnowhere\b", "\bnobody\b", "\bnever\b", and "n't\b".

The highest value for every metric is bolded, whereas the lowest appears in italics. The majority (seven out of 10) of the highest and lowest values are concentrated in Styles A and D. Moreover, each of the metrics differentiates the styles appreciably, with the highest measurements at least double the lowest in all the count-based categories (Negation, Exclamation, and Superlative), and the highest measurements roughly 50% higher than the lowest in the ratio-based categories (Adjective and Adverb).

4.7 Reflection on Automated Metrics

It is crucial to keep in mind that the metrics serve merely as loose proxies for the various stylistic elements discussed in the critical and pedagogical literature. The metrics do not provide absolute judgments for confirming or refuting scholarly claims. Rather, they offer a wholly distinct method of analysis from the kind traditionally undertaken by literary experts. As Lisa M. Rhody

eloquently reasons in support of her application of computational topic modeling techniques to poetry,

> [...] researchers can make use of a powerful tool with which to explore latent patterns in poetic texts. For poetry data in particular and literary texts in general, close reading and contextual understanding work together, like the weaving and unraveling of Penelope at her loom, in order to identify relations between texts by shuttling between computational de-familiarization and scholarly experience. (Rhody, 2012)

Agreement between the "de-familiarizing" computational angle and the more traditional "scholarly" perspective is a valuable site of examination, and so are non-agreement and disagreement. There is much to gain from reasoning about what one method reveals that the other does not, or about what one method proposes that the other contradicts.

Chapter 5
METHODOLOGY

This chapter begins by describing the architecture of the text generation system. It then lays out the two stages of the training procedure. The last section unites all of the hyperparameters introduced in this chapter into three tables for ease of viewing and comparison.

The first stage of the training procedure is necessary for preparing one component of the text generation system for the second stage. During the second stage, five distinct conditions are introduced, and each one is run once, with a random seed for reproducibility. The conditions of the second stage represent different configurations of stylistic influence.

5.1 The Models

From a bird's eye view, the proposed system consists of a text generator (henceforth simply the *generator*) and a style discriminator (the *discriminator*). The discriminator judges how likely a sample sentence is to have emanated from each of the styles that it has been primed to identify. When the discriminator assigns greater probabilities to the styles that the generator is configured to prefer and to emulate, the generator incurs less loss.

Both the generating and discriminating components of the system are warm-started from the historically pre-trained MacBERTh (Manjavacas et al., 2022). The discriminator consists of a MacBERTh checkpoint fitted with a randomly-initialized multi-class classification head and a dropout regularization rate of 40% in the classification layer.

The generator consists of the same MacBERTh checkpoint fitted with a randomly-initialized next-token prediction head that operates over a concatenation of three different views (or representations) of the input: the last hidden state of the last seen token, the max-pooling aggregation of the last hidden states of all seen tokens, and the pooler output derived from the last hidden state of the [CLS] token.[7] These three equally-sized representations ensure that the prediction of the next-token takes into account the full context of the tokens seen so far in addition to the local context, i.e. any probabilistic constraints incurred by the immediately preceding token.

[7] [CLS] is a special token used across the BERT family to indicate the beginning of an input sequence. The final representation of the [CLS] token is generally pre-trained to capture global information and is frequently borrowed as a convenient fixed-length encoding of any variable-length input sequence.

The concatenation of the three representations then passes through the following transformations:

1. Forty percent of the values across the concatenated representations are selected at random (during each forward pass) for dropout.
2. Two feed-forward layers lead to a final output layer with the size of the vocabulary (30,522).
3. If required, Gumbel-Softmax (Jang et al., 2016) is applied to the final output layer to produce a differentiable one-hot tensor, effectively offering a discrete sampling result without fragmenting the computation graph. The Gumbel-Softmax trick allows for backpropagation to flow through the sampling operation.
4. The output of step 2 is returned, along with that of step 3, if required.

For the reader's reference, a PyTorch implementation of the generator's architecture is demonstrated in Appendix A.

5.2 Priming Stage

This stage applies only to the discriminator. The discriminator is trained on the same task that it will perform later on in the adversarial stage: multi-class classification. The classification head plus all layers (including the 12 attention layers) of the base model are permitted to update.

Training samples are drawn from a pool composed of 3584 sentences from each style; separate validation and evaluation sets are constructed each with 256 sentences from each style. Given a batch size of 128, every epoch consists of 140 update steps; inference over the validation set at the end of every epoch requires 10 iterations, and the final pass through the evaluation set at the conclusion of training requires another 10. Training proceeds until the epoch-end validation loss worsens twice consecutively; with a learning rate of 4e-6, this takes a total of 15 epochs, with the 10th epoch's parameters being retained for the adversarial stage.

(Hyperparameters such as batch size and learning rate were selected via grid search with the goal of optimizing for validation loss.)

At the end of the priming stage, the discriminator achieves a macro-averaged and weighted-averaged F1 score of 0.83 on the evaluation set. An exact breakdown of the primed model's classifications and misclassifications is shown in Table 5.2. From here on, the entire discriminator is kept frozen, as it lacks a meaningful loss function in the next stage.

Table 5.2: Style Discriminator Confusion Matrix[8]

True/Pred.	A	B	C	D	E
A	225	15	6	9	1
B	8	219	8	12	9
C	11	7	210	18	10
D	17	17	14	197	11
E	4	10	14	17	211

5.3 Adversarial Stage

5.3.1 The Intrinsic Style

Unlike the discriminator, whose training is distributed evenly across all five literary styles, the generator trains principally on a single pre-selected style – its *intrinsic* style. Style D has been chosen as the generator's intrinsic style across all experiments conducted for this study.[9]

The bread-and-butter of the generator's training regime is teacher-forced next-token prediction over Style D sentences. This is best illustrated via example. Consider the first sentence (tokenized) of "Greenleaf":

> [CLS] mrs . may ' s bed ##ro ##om window was low and faced on the east and the bull , silver ##ed in the moon ##lig ##h ##t , stood under it , his head raised as if he listened - - like some patient god come down to woo her - - for a stir inside the room . [SEP][10] (O'Connor, 1956)

In the above example, white space is used to demarcate token boundaries. At time step 0, the generator receives as input the index for the "[CLS]" token. Based on that information alone, it must attempt to predict the next token in the sequence. (The true next token, of course, is "mrs" without a period.) The generator computes a real numerical value for each token in the vocabulary such that a higher value means a higher probability for the corresponding token; this is the generator's "prediction" in raw form. These numerical values can be transformed into a

[8] Labels along the left margin indicate the true labels; labels across the upper margin indicate the predicted labels.
[9] The choice of intrinsic style is random.
[10] [SEP] is another special token used across the BERT family. It serves to separate sentences when a task involves pairs of sentences; for other tasks, [SEP] can simply indicate the end of a sequence.

probability distribution and then compared to the true distribution (which, in this case, is known to be 1.0 for “mrs” and 0.0 for every other token) to calculate the generator’s prediction error, or *loss*.[11] At no point under this training paradigm is it actually necessary to sample from the predicted probability distribution and pick a next token; calculating the loss alone is sufficient.

Then, regardless of the generator’s performance at time step 0, at time step 1 it is simply shown the “correct answer” (i.e. “mrs”) and – as if assuming that the generator has predicted every token flawlessly up to the present time step – asked again, *what would come next?* (Hence the terms “teacher-forcing” and “teacher-forced.”) Concretely, the generator receives as input the indices corresponding to “[CLS] mrs” and again produces a vocabulary-sized tensor of values to be compared to the true distribution (1.0 for “.” and 0.0 for all other tokens). In this example, a loss would be calculated and saved for each of the rows in Table 5.3.1:

Table 5.3.1: Next-Token Prediction Illustration

Timestep	**Input Sequence**	**Target**
0	[CLS]	mrs
1	[CLS] mrs	.
2	[CLS] mrs .	may
3	[CLS] mrs . may	’
…	…	…
60	[CLS] mrs . may ’ s bed ##ro ##om window was low and faced on the east and the bull , silver ##ed in the moon ##lig ##h ##t , stood under it , his head raised as if he listened - - like some patient god come down to woo her - - for a stir inside the room .	[SEP]

Finally, the accumulated losses are scaled to give higher weight to later time steps and backpropagated all together through the model.

To train the generator on its intrinsic style, 4360 sentences are randomly selected from Style D and shuffled into batches of 20 sentences, for a total of 218 update steps per epoch. The

[11] specifically cross entropy loss

embedding layer and lower eight attention layers of the base model are kept frozen;[12] all other weights, including those of the randomly initialized next-token prediction head, update with a learning rate of 5e-5.

(Unlike in the priming stage, hyperparameter selection in the adversarial stage is primarily driven by memory constraints – for example, in the case of layer freezing and batch size – and by human monitoring of system outputs. Failure states like mode collapse are visually apparent in generated samples; training stability can be ascertained via inspection of loss logs. The final configuration is one that results in sensible and varied generation rather than one that optimizes for a quantitatively precise objective.)

Differences in the training regime arise later on between the control and experimental conditions; however, at the start of the adversarial stage, the generator trains according to the described procedure for 15 epochs, regardless of the condition to which it is eventually assigned.

5.3.2 The Influencing Style

At the end of the initial 15 epochs, the teacher-forced training regime outlined in Section 5.3.1 is perturbed by the introduction of training signal from the discriminator. The teacher-forced regime continues as before but is no longer the sole source of training signal for the generator. Training signal from the discriminator is precisely the lever for differentiating the control and experimental conditions, which are described next.

In order to fulfill its motive of establishing proof-of-principle, this study explores a limited set of stylistic influence configurations. Specifically, it looks at the effect of imposing a single influencing style to perturb an otherwise stylistically homogeneous training process. The pre-selected intrinsic style, as a reminder, is Style D, whereas the influencing styles are the set of all available styles. The conditions will henceforth be denoted *DA*, *DB*, *DC*, *DD*, and *DE*, where the first letter indicates the intrinsic style and the second letter the influencing style. DD serves as the control condition, whereas DA, DB, DC, and DE are experimental conditions.

Starting with epoch 16, every update involves the following steps:

1. Loss from teacher-forced training over the ingested batch is computed.

[12] This choice is a consequence of memory constraints. In particular, an inverse relationship exists between the number of weights that are allowed to update and the batch size.

2. The generator now auto-regressively[13] generates new sequences. The number of sequences it generates is defined by the same batch size as in step 1. The new sequences are constructed via sampling from the generator's output distributions; since there will be a loss associated with these new sequences (see step 4), *the sampling operation must be differentiable.* Given these requirements, *the generator must pass its output through Gumbel Softmax before returning it.*
3. The generated sequences are judged by the discriminator. This results in another loss term, which will be called the *adversarial loss*.
4. After some scaling, the two losses are combined and backpropagated together through the generator.

Step 3 merits additional explanation. The adversarial loss is designed to encourage the generator to produce sequences similar to those of its intrinsic style *or* of the influencing style; both are treated as targets. Thus, given a condition D*X* and some generated sample, the adversarial loss is calculated as the negative log classification probability (as computed by the discriminator) of the sample belonging to the intrinsic style (Style D), plus the negative log classification probability of the sample belonging to Style *X*. (Note: When *X* is Style D, i.e. under the control condition, the adversarial loss is made up simply of the negative log classification probability of the sample belonging to the intrinsic style.)

In other words, the discriminator's predicted distribution is evaluated against a multi-label target, where both the influencing style *X* and the intrinsic style D are coded as correct, and all other styles as incorrect. In effect, this multi-label target incentivizes the generator to produce outputs that lead to the discriminator amassing as much probability as possible on the "correct" classes (or *class* in the singular, under the control condition).

As outlined in step 4, the adversarial loss is scaled and added to the teacher-forced loss and backpropagated together through the generator. In practice, penalizing the generator for producing outputs that are assigned high probabilities for the *incorrect* classes was detrimental to model performance and often induced mode collapse (a frequently observed failure mode in generative adversarial training). Therefore, all losses related to incorrect classes were zeroed out,

[13] Under the teacher-forcing paradigm, sequences are incremented using the known "correct" tokens. In contrast, in the auto-regressive setting, the generator is entirely self-sufficient. At each time step, the generator samples from its predicted distribution and then appends the sampled token to the end of the input sequence, thus forming the input sequence for the next time step. It continues like this until it produces an [SEP] token.

and only losses incurred from the correct classes were conserved. It was also found that, when combining the teacher-forced and adversarial losses, scaling the latter by a positive factor less than 1 (0.8 was used in this study) was beneficial for training stability.

Training under this joint loss continues for an additional 40 epochs, i.e. epochs 16-55. All other configurations, such as weight freezing and learning rate, are unchanged from the first 15 epochs. At the conclusion of every epoch, 500 sample sequences are elicited from the generator.

5.4 Summary of Configurations

5.4.1 Model Hyperparameters

Table 5.4.1 summarizes the hyperparameters pertaining to the generator and the discriminator.

Table 5.4.1: Model Hyperparameters

Hyperparameter	**Generator**	**Discriminator**
Base Model	MacBERTh	MacBERTh
Base Model Default Hyperparameters	Refer to Manjavacas et al., 2022	Refer to Manjavacas et al., 2022
Base Model Maximum Sequence Length	64	64
Next-Token Prediction Head	Refer to Appendix A for full schematic	N/A
Prediction Head Input Layer	Concatenated representations from Base Model	N/A
Prediction Head Input Layer Size	2304	N/A
Prediction Head Input Layer Dropout Regularization	0.4	N/A
Prediction Head Fully-Connected Layer Size	324	N/A
Prediction Head Fully-Connected Layer	Leaky ReLU with slope of 0.1	N/A

Activation		
Prediction Head Output Layer Size	30,522	N/A
Gumbel-Softmax Temperature (Tau)	2.5	N/A
Classification Head	N/A	BertForSequenceClassification
Classification Head Default Parameters	N/A	Refer to BertForSequenceClassification module from Huggingface Transformers library
Number of Classes	N/A	5
Dropout Regularization	N/A	0.4

5.4.2 Training Hyperparameters

Tables 5.4.2A and 5.4.2B summarize the hyperparameter settings pertaining to the training loops of the priming and adversarial stages.

Table 5.4.2A: Priming Stage Configuration

Configuration Element	**Value for Priming Stage**
Task	Text Classification
Trained Models	Discriminator
Frozen Layers	None
Other Involved Models	N/A
Training Dataset	3584 randomly sampled sentences per style
Validation Dataset	256 randomly sampled sentences per style
Evaluation Dataset	256 randomly sampled sentences per style
Dataset Total	4096 randomly sampled sentences per style
Batch Size	128
Maximum Number of Epochs	30 (in practice, training stopped after 15 epochs,

	with epoch 10 yielding best model)
Early Stopping Condition	Validation loss worsens (increases) two epochs in a row
Loss Function	Cross Entropy Loss (computed within BertForSequenceClassification module)
Optimizer	AdamW from Pytorch optimizers
Learning Rate	4e-6
Random Seed	15324

Table 5.4.2B: Adversarial Stage Configuration

Configuration Element	**Value During First 15 Epochs of Adversarial Stage**	**Value During Last 40 Epochs of Adversarial Stage**
Task	Teacher-forced next-token prediction	Teacher-forced next-token prediction AND Adversarial generation
Trained Models	Generator	Generator
Frozen Layers	macberth.embeddings, macberth.encoder.layer.0, macberth.encoder.layer.1, macberth.encoder.layer.2, macberth.encoder.layer.3, macberth.encoder.layer.4, macberth.encoder.layer.5, macberth.encoder.layer.6, macberth.encoder.layer.7	macberth.embeddings, macberth.encoder.layer.0, macberth.encoder.layer.1, macberth.encoder.layer.2, macberth.encoder.layer.3, macberth.encoder.layer.4, macberth.encoder.layer.5, macberth.encoder.layer.6, macberth.encoder.layer.7
Other Involved Models	N/A	Primed Discriminator
Training Dataset	4360 randomly sampled sentences from Style D	4360 randomly sampled sentences from Style D (*teacher-forced next-token prediction task*)
Validation Dataset	N/A	N/A
Evaluation Dataset	N/A	N/A

Dataset Total	4360 randomly sampled sentences from Style D	4360 randomly sampled sentences from Style D (*teacher-forced next-token prediction task*)
Batch Size	20	20 (*teacher-forced next-token prediction task*) AND 20 (*adversarial generation task*)
Maximum Number of Epochs	15	40
Early Stopping Condition	None, all epochs must be completed	None, all epochs must be completed
Loss Function	Cross Entropy Loss with boosting of infrequent timesteps[14]	Cross Entropy Loss with boosting of infrequent timesteps (*teacher-forced next-token prediction loss*) AND Scaled[15] Weighted[16] Cross Entropy Loss (*adversarial loss*)
Optimizer	AdamW from Pytorch optimizers	AdamW from Pytorch optimizers
Learning Rate	5e-5	5e-5
Random Seed	15324	15324

[14] This is implemented as cross_entropy_loss = cross_entropy_loss / (number_of_non-padded_targets_at_timestep_t - 0.5), where number_of_non-padded_targets_at_timestep_t is at maximum equal to the batch size and at minimum equal to 1. The described transformation of cross_entropy_loss increases the weight placed on losses incurred at later timesteps where the training data generally becomes more scarce.

[15] The scaling coefficient for the adversarial loss is 0.8.

[16] The weighting vector is defined as [*A*, *B*, *C*, 1, *E*], where *A*, *B*, *C*, and *E* are all indicator functions that are true when the name of the influencing style bears their letter and false otherwise. The adversarial loss related to uninvolved styles is thereby zeroed out.

Chapter 6
RESULTS AND DISCUSSION

The previous chapter laid out five influence configurations, all with Style D as their intrinsic style and varying in their specific influencing style (or lack thereof). This chapter brings back the 10 style metrics introduced in Chapter 4 and applies them to the outputs generated under the five conditions, in order to compare the generative behaviors resulting from those conditions.

More precisely, for each condition, the metrics are computed over the 500 final sample sentences, i.e. those elicited from the generator upon the completion of the 40th epoch of joint loss training (the 55th epoch in total). Results are collected in Table 6.

The following specifications apply:

- All metrics take the entirety of the 500 sentences as their basis. There is no random sampling performed over these sentences.
- The formulae for calculating the metrics are exactly as described in Chapter 5.
- The only difference is found in the sample sizes, upon which the formulae do not vary. The only metric for which the expected value is sensitive to sample size is *Type-Token Ratio*. (In theory, values should decrease as sample size increases and, given any fixed size vocabulary, approach a limit of 0.)
- *Type-Token Ratio* uses a random sample of 6000 tokens drawn from the 500 sentences.
- *Percentage of 3+ Syllable Tokens* uses the same random sample of 6000 tokens as *Type-token Ratio*.
- *Average Syllable Count per Single-Breath Segment* uses a random sample of 500 segments drawn from the 500 sentences.

Table 6: Tabulated Results[17]

Metric	**DA**	**DB**	**DC**	**DD**	**DE**
Type-Token Ratio	0.284	0.265	0.286	0.268	0.297

[17] calculated over 500 elicited generations at the completion of the 55th epoch (control condition in red)

% of 3+ Syllable Tokens	8.5	7.6	8.74	7.34	8.44
Average Syllable Count per Single-Breath Segment	11.836	11.844	11.314	11.08	11.926
% of Straight-Through Sentences	29.4	33.6	31.	31.8	34.4
% of Sentences Containing Profanity	3.	2.6	2.6	2.2	3.
Negation	30.	29.8	30.4	30.4	28.2
Adjective	0.167	0.171	0.181	0.158	0.172
Adverb	0.337	0.332	0.342	0.312	0.35
Exclamation	5.2	6.4	6.6	5.8	4.
Superlative	0.8	1.4	1.2	0.	0.6

Additionally, a line plot is produced for every metric and included in Appendix B. Each line plot shows one metric computed for the five conditions throughout the course of the joint loss training; the points represent the metric values calculated for 500 sample sentences elicited at regular intervals (epoch 15, 25, 35, 45, and 55).

Note that the values across all five conditions have the same starting position, except where the computation of the metric involves random sampling (i.e. Type-Token Ratio, % of 3+ Syllable Tokens, and Average Syllable Count per Single-Breath Segment). The starting positions coincide because there is no difference between the conditions during the first 15 epochs of training (teacher-forced next-token prediction only), and moreover since the shuffling of training batches is determined by a random seed. This equivalence helps provide a clearer picture of how the presence of the adversarial loss affects the evolution of the metrics' values over time.

Note also that the plots suggest a moderate oscillation shared across all conditions induced by the joint training signal.

Finally, for the sake of transparency and interpretability, Appendix C shows sample system outputs, as decoded by the same tokenizer that was used previously for tokenizing the input sentences. For each condition, the same fixed selection of 50 (out of 500 total) outputs from the final epoch are displayed. Here again, the benefits of the random seed are clear. Samples elicited at the end of any given epoch are directly comparable across the conditions; they are often identical and generally deviate only at specific junctions in the generated sequences. This direct comparability facilitates a potential close reading of the system outputs as well as other modes of qualitative analysis.

6.1 General Analysis

As reasoned in Chapter 5, each of the 10 metrics sheds light on a certain aspect of literary style. Through analysis of the experimental results, the following question will be addressed: *For each metric, does the aspect of style encapsulated by that metric shift as hypothesized?*

Toward this end, the experimental conditions are each compared to the control condition. For any given metric, the direction and magnitude of their difference reveals how the system responded to the influence configuration encoded by the particular experimental condition.

Some aspects of style shift more convincingly than others; some aspects do not appear to be acted upon at all. Specifically:

- There is weak support for the conclusion that the aspect represented by the *Type-Token Ratio* metric shifted as intended. It is discussed in greater detail in the error analysis section.
- Based on the initial analysis of corpus styles, conditions DC and DA are expected to generate the greatest and second greatest increases in the *Percentage of 3+ Syllable Tokens* metric, respectively. Condition DB is expected to raise the metric lightly. These predictions are borne out. Condition DE, on the other hand, presents a minor anomaly. Overall, the intended shifts in the aspect represented by this metric enjoy moderate support.

- For the aspect reflected by the *Average Syllable Count per Single-Breath Segment* metric, results contain no evidence that the intended shift occurred. It is discussed in the error analysis section.
- Likewise, for the aspect reflected by the *Percentage of Straight-Through Sentences* metric, no evidence was found for the intended shift. It is discussed in the error analysis section.
- Regarding the *Profanity* metric, conditions DE, DB, and DA are expected to exert a numerically positive (and pedagogically negative) influence on the impressionable generator. There is evidence that the generator evolves to become more profane under those conditions than in the control condition. However, the profanity value of condition DC also rests higher than under the control condition; the opposite is predicted. Overall, in light of the small margin between the lowest and highest profanity values (0.8, or four sentences out of 500), the results are interpreted as offering weak support for the hypothesized shifts on this aspect of style.
- The influence of other styles in the experimental conditions can be observed generally pulling the *Negation* metric downward, compared to the control condition; this is as expected. However, condition DB is expected to pull the negation metric higher, since Style B has a significantly higher negation metric than the other styles; the data does not reflect this. Overall, for the aspect of style approximated by the Negation metric, there is weak support for the hypothesized shift.
- The experimental conditions can be observed pulling both the *Adjective* and the *Adverb* metrics higher. This is predicted by the analysis of initial styles, which revealed that Style D has the lowest average adjective to noun and adverb to verb ratios in the corpus. In particular, conditions DE and DC exert the greatest upward pressure on the adverb metric, which is also predicted by the initial analysis. The results lend moderate and strong support, respectively, for the hypothesized shifts in the aspects represented by the adjective and adverb metrics.
- See the error analysis section for discussion of the *Exclamation* metric result. The aspect represented by this metric does not appear to undergo the intended style shift.
- In spite of possibly weak signal from the low occurrence of superlatives across the corpus, the expected trend of experimental conditions pulling the *Superlative* metric

higher is indeed observed. The control condition produces no superlatives in the final output of 500 sentences, whereas between three and seven superlatives are produced under the experimental conditions. There is moderate support for the intended style shift pertaining to this metric.

6.2 Error Analysis

6.2.1 Type-Token Ratio

The Type-Token Ratios of the corpus styles suggest that influence from Styles A and C should lead to higher ratios, Style E should lead to a lower ratio, and Style B should have little effect relative to the control condition. Conditions DA, DB, and DC do in fact bear out these predictions. Unexpectedly, however, condition DE results in the highest Type-Token Ratio of all the conditions.

Why might this be the case? Vocabulary is one possible explanation. If Style E exhibits a distribution of tokens that is most different from that of Style D, the DE condition could well produce a higher-than-expected type-token ratio as the generator attempts to accommodate both Styles D and E.

A back-of-the-envelope calculation of cosine similarities confirms the plausibility of this hunch. The same 24,000 tokens sampled for each style in the initial computation of corpus metrics (refer to Chapter 5) can be re-used to generate an estimate of each style's token distribution. Then, each style's token distribution can be compared to that of Style D using a measure like cosine similarity.

Table 6.2.1 shows the results from following exactly that process. As suspected, in terms of token distribution, Style E is substantially *less similar* to Style D than any of the other styles are. This could explain the anomalous result of the DE condition; if so, Type-Token Ratio might not be the most appropriate metric for stylistic influence experiments, since the similarity between token distributions would constitute a powerful confounding factor.

Table 6.2.1: Token Distribution Cosine Similarities Between Corpus Styles

Metric	A-D	B-D	C-D	D-D	E-D
Cosine Similarity	0.947	0.964	0.907	1.	0.819

between Token Distributions					

6.2.2 Average Syllable Count per Single-Breath Segment

Under all conditions, the metric fell into the range of 11 to 12 syllables. The lack of adaptation to influencing styles according to this metric suggests that the system failed to pick up on the pause tendencies of the corpus styles or, if it did learn those patterns, failed to apply that knowledge in its generation process.

Compared to token-level metrics like Type-Token Ratio, Profanity, and Superlative, this metric resides on a much higher level of abstraction. Since the generator is a next-token predictor that, naturally, operates on the token level, higher levels of abstraction might not be learned without explicit emphasis.

6.2.3 Percentage of Straight-Through Sentences

Just like the previous metric, the Percentage of Straight-Through Sentences metric is relatively abstract. Explicit representation of syntax (or of the heuristics that the metric relies on to determine “straight-through-ness”) might draw more attention to this construct and facilitate the system’s learning thereof.

6.2.4 Exclamation

Given that exclamation marks appear only up to 5 times per 100 sentences in any of the corpus styles, it is probable that the discriminator failed to pick up on any association between exclamation mark frequency and literary style – and that the signal it passed on to the generator was resultantly weak. In general, this style metric’s narrow lens might limit its utility.

6.3 Summary of Results

To recapitulate, the methodology proposed by this study appears to be most conducive toward inducing style shifts measurable by the following metrics:

- Percentage of 3+ Syllable Tokens
- Adjective
- Adverb

- Superlative

Style shifts in the aspects tracked by the following metrics are also observed, albeit weakly:

- Type-Token Ratio
- Profanity
- Negation

Finally, for reasons explored in the error analyses, the aspects of style measured by the following metrics do not respond to the proposed methodology:

- Average Syllable Count per Single-Breath Segment
- Percentage of Straight-Through Sentences
- Exclamation

Specifically, the generator is unlikely to pick up on style aspects that are poorly represented in the corpus (a sparsity problem) or that involve greater removal from the basic building block of the generative process, i.e. the token (an abstraction problem).

Chapter 7
CONCLUSION

"The isolated man does not develop any intellectual power," Turing wrote in 1948. Only through sustained interaction with a social environment might an individual "make a very few discoveries which are passed on […]" (Turing, 1948).

Born out of an observation expressed by Turing, Boden, and others that creativity does not happen in a vacuum, this study seeks to pass along its own "few discoveries." Concretely, it conceives, implements, and evaluates an approach to modeling stylistic influence in text generation. In accomplishing the above, it provides a proof-of-principle for leveraging adversarial signal as a way to induce stylistic shift – but more so along certain dimensions than others. Dimensions most amenable to the proposed paradigm are characterized by a low level of abstraction with respect to individual tokens and by sufficient presence in the training data. These dimensions include diction (as measured by the Percentage of 3+ Syllable Tokens metric) and descriptiveness (as measured by the Adjective and Adverb metrics).

On the other hand, style shift along dimensions like rhythm (as measured by the Average Syllable Count per Single-Breath Segment metric) and sentence complexity (as measured by the Percentage of Straight-Through Sentences metric) proves elusive. Moreover, while Type-Token Ratio is indeed a legitimate and useful metric for quantifying diction from a vocabulary richness angle, it might not be an appropriate metric for quantifying style evolution, because how much this metric changes depends also on how similar the intrinsic and influencing styles are to each other.

7.1 Limitations

The language modeling power of the implemented system is limited first and foremost by the size of the 1950s Anglophone Styles of Fiction corpus. At only 4096 sentences for the style with the fewest exemplars and 4169 sentences for the intrinsic style selected for this study, the corpus can quickly be overfit by a composite generator-discriminator system with a combined count of over 200 million learned parameters. In practice, prolonged training was found to encourage memorization of the training data. Thus, addressing the size of the corpus must come before entertaining an increase in the size of the models.

Furthermore, decisions involving numerous hyperparameters – such as batch size, pre-trained parameter freezing, and maximum sequence length – were necessarily tied to memory constraints of the available computing resources. As a result of the low ceiling on sequence lengths, this study was restricted to the generation and evaluation of stand-alone sentences, which are sufficient in theory but not so much in practice.

7.2 Future Research

The present study opens onto numerous avenues of continued research. Four of them are highlighted below.

Having observed the failure of the implemented system to induce shifts along the more abstract dimensions of style, one might explore mechanisms for nudging the system toward placing greater attention on those specific dimensions. Possible mechanisms include injecting features that explicitly describe a particular dimension into model inputs or integrating additional system components that preside over the dimension.

Second, the pantheon of stylometry remains far from full: The exploration of 10 style metrics undertaken in the present study is but one more foray into the space of stylometry techniques. There exists vast potential for the development of style metrics that deviate from or build upon the standard procedure of leveraging stop words.

Third, this study illuminates the effects of adding just a single positive peer influence. But what would it look like to incorporate multiple positive influences? Or a negative, repelling influence? Or a combination of both? In short, this study can be expanded to probe the effects of more naturalistic influence configurations.

And lastly, the system demonstrated in this study finds itself on the doorstep of grander possibilities in social creativity research. Armed with the flexibility to represent nuanced and dynamic influence profiles, this system would make an ideal subject in an Artificial Creative System (Saunders, 2012), such as a clique formation simulation (Zhang, 2017), or other practical explorations of Csikszentmihalyi's Domain-Individual-Field-Interaction paradigm (Csikszentmihalyi, 1999).

A literary soirée of text generation systems, pushing and pulling at one another, and then pushing and pulling some more, might finally find their way up to "the brightest heavens of invention" (Shakespeare et al., 1883).

BIBLIOGRAPHY

Achebe, C. (1950). Polar Undergraduate. *University Herald*, University of Ibadan.

Achebe, C. (1951). In a Village Church. *University Herald*, University of Ibadan.

Achebe, C. (1952). Marriage is a Private Affair. *University Herald*, University of Ibadan.

Achebe, C. (1953). Dead Men's Path. *University Herald*. University of Ibadan.

Achebe, C. (1958). *Things Fall Apart*. London: Heinemann.

Achebe, C. (1959). The Sacrificial Egg. *The Atlantic*.

Achebe, C. (1965a). English and the African Writer. *Transition*, 18. https://doi.org/10.2307/2935429.

Achebe, C. (1965b). The Novelist as Teacher. *New Statesman*, London.

Als, H. (1998). The Enemy Within: The making and unmaking of James Baldwin. *The New Yorker.*

Armstrong, J. (2001). Blinded by Whiteness: Revisiting Flannery O'Connor and Race. *Flannery O'Connor Review*, *1*, 77–86. http://www.jstor.org/stable/26669737.

Ashton, H. (2018). "I'll Come Back and Break Your Spell": Narrative Freedom and Genre in *The Haunting of Hill House*. *Style*, *52*(3), 268–286. https://doi.org/10.5325/style.52.3.0268.

Baldwin, J. (1953). *Go Tell It on the Mountain*. New York, NY: Alfred A. Knopf.

Baldwin, J. (1957). Sonny's Blues. *Partisan Review*.

Bird, S.; Klein, E.; and Loper, E. (2009). *Natural Language Processing with Python*. O'Reilly Media, Inc.

Boden, M. A. (2003). *The Creative Mind: Myths and Mechanisms*, 2nd Edition. New York, NY: Routledge.

Boden, M. A. (2009). Creativity in a Nutshell. *Think* 5:83–96.

Burns, S. L. (1974). O'Connor and the Critics: An Overview [Review of *The Eternal Crossroads: The Art of Flannery O'Connor; The Christian Humanism of Flannery O'Connor; Nightmares and Visions; Invisible Parade: The Fiction of Flannery O'Connor; The Question of Flannery O'Connor; Flannery O'Connor*, by L. V. Driskell, J. T. Brittain, D. Eggenschwiler, G. H. Muller, M. Orvell, M. Stephens, & D. Walters]. *The Mississippi Quarterly*, *27*(4), 483–495. http://www.jstor.org/stable/26474262.

Burrows, J. (2002). Delta: A Measure of Stylistic Difference and a Guide to Likely Authorship. *Literary and Linguistic Computing*, 17 (3), 267–287.

Chang, K. K.; Cramer, M.; Soni, S.; and Bamman, D. (2023). Speak, Memory: An Archaeology of Books Known to ChatGPT/GPT-4. *ArXiv*, 2305.00118.

Colombo, G. B.; Liu, H.; and Whitaker, R. M. (2023). Exploring Human Models of Innovation for Generative AI. In *Proceedings of the 14th International Conference on Computational Creativity*.

Costello, D. P. (1959). The Language of "The Catcher in the Rye." *American Speech*, *34*(3), 172–181. https://doi.org/10.2307/454038.

Costello, D. P. (1963). Salinger & His Critics: Autopsy of a Faded Romance. *Commonweal*.

Csikszentmihalyi, M. (1999). *A Systems Perspective on Creativity*. Cambridge, UK: Cambridge University Press.

Danescu-Niculescu-Mizil, C.; Lee, L.; Pang, B.; and Kleinberg, J. (2012). Echoes of Power: Language Effects and Power Differences in Social Interaction. In *Proceedings of WWW*.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. (2019). BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In Burstein, J.; Doran, C.; and Solorio, T., eds., *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Volume 1 (Long and Short Papers), 4171–4186. Minneapolis, Minnesota: Association for Computational Linguistics.

D'Souza, L. and Mimno, D. (2023). The Chatbot and the Canon: Poetry Memorization in LLMs. In *CHR 2023 CEUR Workshop Proceedings*.

Ellsberg, R. (2005). Flannery O'Connor: Spiritual Master. *American Catholic Studies*, *116*(1), 59–72. http://www.jstor.org/stable/44194864.

Francis, C. (2014). *The Critical Reception of James Baldwin, 1963-2010: "An Honest Man and a Good Writer"*. Boydell & Brewer, Camden House.

Franklin, R. (2008). After Empire: Chinua Achebe and the great African novel. *The New Yorker*.

Friedman, M. J. (1962). Flannery O'Connor: Another Legend in Southern Fiction. *The English Journal*, *51*(4), 233–243. https://doi.org/10.2307/810724.

Gemini Team et al. (2023). Gemini: A Family of Highly Capable Multimodal Models. *arXiv*, 2312.11805.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. (2014). Generative Adversarial Networks. *Advances in Neural Information Processing Systems* 3.

Hague, A. (2005). "A Faithful Anatomy of Our Times": Reassessing Shirley Jackson. *Frontiers: A Journal of Women Studies*, *26*(2), 73–96. http://www.jstor.org/stable/4137397.

Hawkes, J. (1962). Flannery O'Connor's Devil. *The Sewanee Review*, *70*(3), 395–407. http://www.jstor.org/stable/27540793.

Hosseini, K.; Beelen, K.; Colavizza, G.; and Coll Ardanuy, M. (2021). Neural Language Models for Nineteenth-Century English. *Journal of Open Humanities Data* 7.

Jackson, S. (1959). *The Haunting of Hill House*. New York, NY: Viking Press.

Jang, E.; Gu, S. S.; and Poole, B. (2016). Categorical Reparameterization with Gumbel-Softmax. *ArXiv*, abs/1611.01144.

Jordanous, A. (2012). A Standardised Procedure for Evaluating Creative Systems: Computational Creativity Evaluation Based on What It Is to Be Creative. *Cognitive Computation* 4.

Lamb, C.; Brown, D. G.; and Clarke, C. L. A. (2018). Evaluating Computational Creativity: An Interdisciplinary Tutorial. *ACM Computing Surveys* 51(2).

Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. (2019). RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv*, abs/1907.11692.

Luke, S. (2023). Computational Creativity as Dynamic, Multiobjective, Multiagent Optimization. In *Proceedings of the 14th International Conference on Computational Creativity*.

Manjavacas, E., and Fonteyn, L. (2022). Adapting vs. Pretraining Language Models for Historical Languages. *Journal of Data Mining Digital Humanities NLP4DH*.

McCarthy, B. E. (1985). Rhythm and Narrative Method in Achebe's "Things Fall Apart." *NOVEL: A Forum on Fiction*, *18*(3), 243–256. https://doi.org/10.2307/1345790.

Muller, G. H. (1972). *Nightmares and Visions: Flannery O'Connor and the Catholic Grotesque*. University of Georgia Press.

O'Connor, F. (1952). *Wiseblood*. New York, NY: Harcourt, Brace and Company.

O'Connor, F. (1955). A Good Man Is Hard to Find. *A Good Man Is Hard to Find and Other Stories*. Boston, MA: Harcourt.

O'Connor, F. (1956). Greenleaf. *The Kenyon Review*.

Ohmann, C., & Ohmann, R. (1976). Reviewers, Critics, and "The Catcher in the Rye." *Critical Inquiry*, *3*(1), 15–37. http://www.jstor.org/stable/1342870.

OpenAI et al. (2023). GPT-4 Technical Report. *arXiv*, 2303.08774.

Parks, J. G. (1984). Chambers of Yearning: Shirley Jackson's Use of the Gothic. *Twentieth Century Literature*, *30*(1), 15–29. https://doi.org/10.2307/441187.

Pease, A.; Colton, S.; and Banar, B. (2023). On the Notion of Creative Personhood. In *Proceedings of the 14th International Conference on Computational Creativity*.

Pennebaker, J. W.; Francis, M. E.; and Booth, R. J. (2001). Linguistic Inquiry and Word Count: A computerized text analysis program.

Pennebaker, J. W. and King, L. A. (2000). Linguistic styles: Language Use as an Individual Difference. *Journal of Personality and Social Psychology*. 77. 1296-312. 10.1037//0022-3514.77.6.1296.

Quayson, A. (1994). Realism, Criticism, and the Disguises of Both: A Reading of Chinua Achebe's "Things Fall Apart" with an Evaluation of the Criticism Relating to It. *Research in African Literatures*, Vol. 25, No. 4 (Winter, 1994), pp. 117-136.

Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. (2019). Language Models Are Unsupervised Multitask Learners. *OpenAI Blog*, 1.

Rhodes, M. (1961). An Analysis of Creativity. *The Phi Delta Kappan* 42(7):305–310.

Rhody, L. M. (2012). Topic Modeling and Figurative Language. *Journal of Digital Humanities*, Vol. 2 No. 1 Winter.

Ritchie, G. (2001). Assessing Creativity. In *Proceedings of the AISB*.

Roemer, D. M. (1992). The Personal Narrative and Salinger's "Catcher in the Rye." *Western Folklore*, *51*(1), 5–10. https://doi.org/10.2307/1499640.

Rosen, G. (1977). A Retrospective Look at the Catcher in the Rye. *American Quarterly*, *29*(5), 547–562. https://doi.org/10.2307/2712573.

Rothe, S.; Narayan, S.; and Severyn, A. (2020). Leveraging Pre-Trained Checkpoints for Sequence Generation Tasks. *Transactions of the Association for Computational Linguistics* 8:264–280.

Rubenstein, R. (1996). House Mothers and Haunted Daughters: Shirley Jackson and Female Gothic. *Tulsa Studies in Women's Literature*, *15*(2), 309–331. https://doi.org/10.2307/464139.

Salinger, J. (1951). *The Catcher in the Rye*. Boston, MA: Little, Brown and Company.

Saunders, R. (2012). Towards Autonomous Creative Systems: A Computational Approach. *Cognitive Computation* 4.

Shakespeare, W., and Kellogg, B. (1883). Shakespeare's King Henry V. *English Classics with Explanatory Notes*. Clark & Maynard.

Sterman, S.; Huang, E.; Liu, V.; and Paulos, E. (2020). Interacting with Literary Style Through Computational Tools. *CHI'20*.

Strunk, Jr., W. (1920). *The Elements of Style*. Harcourt, Boston, MA.

Strunk, Jr., W., and White, E. (1959). *The Elements of Style*. MacMillan, New York, NY.

Torchia, J. (1996). Inside Flannery O'Connor. *The Flannery O'Connor Bulletin*, *25*, 81–102. http://www.jstor.org/stable/26669949.

Turing, A. (2004). Intelligent Machinery (1948). In Copeland, B. J., eds., *The Essential Turing*. Oxford.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.; Kaiser, Ł; and Polosukhin, I. (2017). Attention Is All You Need. *NIPS* 2017.

Walsh, M. (2018). Tweets of a Native Son: The Quotation and Recirculation of James Baldwin from Black Power to #BlackLivesMatter. *American Quarterly*, 70, no. 3.

Wegmann, A.; Schraagen, M.; and Nguyen, D. (2022). Same Author or Just Same Topic? Towards Content-Independent Style Representations. In Gella, S.; He, H.; Majumder, B. P.; Can, B.; Giunchiglia, E.; Cahyawijaya, S.; Min, S.; Mozes, M.; Li, X. L.; Augenstein, I.; Rogers, A.; Cho, K.; Grefenstette, E.; Rimell, L.; and Dyer, C., eds., *Proceedings of the 7th Workshop on Representation Learning for NLP*, 249-268. Dublin, Ireland: Association for Computational Linguistics.

Whitfield, S. J. (1997). Cherished and Cursed: Toward a Social History of The Catcher in the Rye. *The New England Quarterly*, *70*(4), 567–600. https://doi.org/10.2307/366646.

Zhang, A. (2017). A Computational Model of Creative Design as a Sociocultural Process Involving the Evolution of Language. Ph.D. Dissertation, University of Sydney.

APPENDIX A: IMPLEMENTATION OF CUSTOM TEXT GENERATOR

```
import torch
from transformers import BertModel


class CustomGenerator(torch.nn.Module):
    """
    Generator class that stacks token prediction head on top of pre-trained MacBERTh
    Input dims: batch size * sequence length
    Output dims: batch size * vocab size
    """

    def __init__(self, vocab_size):
        super().__init__()
        self.macberth = BertModel.from_pretrained('emanjavacas/MacBERTh')
        self.model_size = self.macberth.config.hidden_size
        self.prediction_head = TokenPredictionHead(self.model_size, vocab_size)

    def forward(self, input_ids, gumbel_tau=2.5, requires_gumbel_out=False):
        macberth_out = self.macberth(input_ids=input_ids)

        last_token_rep = macberth_out.last_hidden_state[:, -1, :]
        max_pooling, _ = torch.max(macberth_out.last_hidden_state, dim=1)
        pooled_output = macberth_out.pooler_output

        prediction_input = torch.cat((last_token_rep, max_pooling, pooled_output), dim=1)
        return self.prediction_head(prediction_input, gumbel_tau, requires_gumbel_out)


class TokenPredictionHead(torch.nn.Module):
    """
    Custom token prediction head that ingests a representation of preceding sequence
    from MacBERTh and passes it through two FF layers
    Input dims: batch size * three times model size
    Output dims: batch size * vocab size
    """
    def __init__(self, model_size, vocab_size):
        super().__init__()
        self.dropout = torch.nn.Dropout(0.4)
        self.fc1 = torch.nn.Linear(model_size * 3, 324)
        self.leaky_relu = torch.nn.LeakyReLU(0.1)
        self.fc2 = torch.nn.Linear(324, vocab_size)
```

```
    def forward(self, x, tau, gumbel_out):
        x = self.dropout(x)
        x = self.leaky_relu(self.fc1(x))
        x = self.fc2(x)

        if gumbel_out:
            one_hot_x = torch.nn.functional.gumbel_softmax(x, tau=tau, hard=True)
        else:
            one_hot_x = None
        return x, one_hot_x
```

APPENDIX B: LINE PLOTS OF METRICS OVER TRAINING EPOCHS

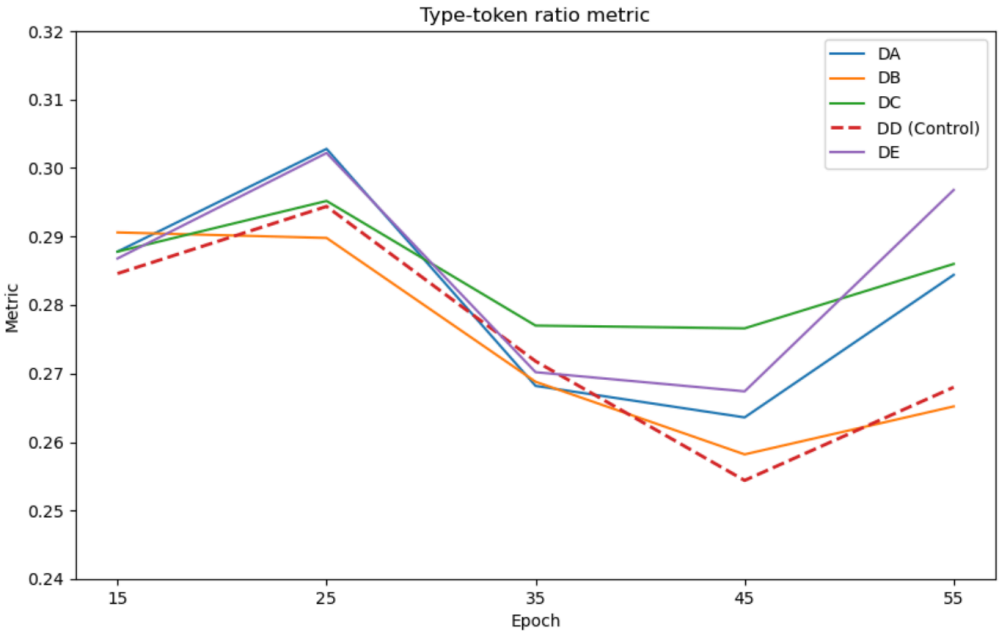

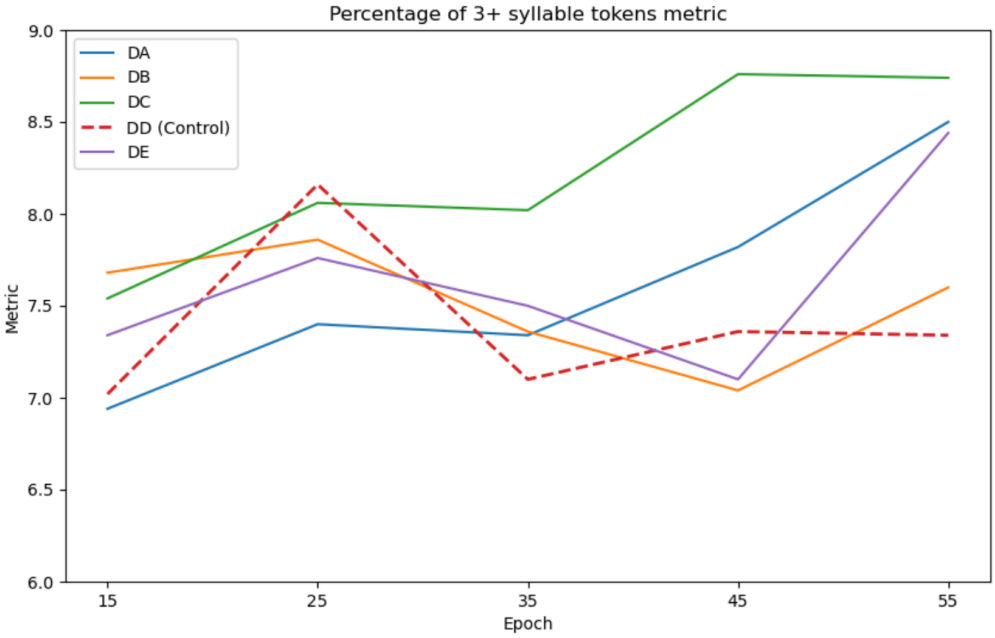

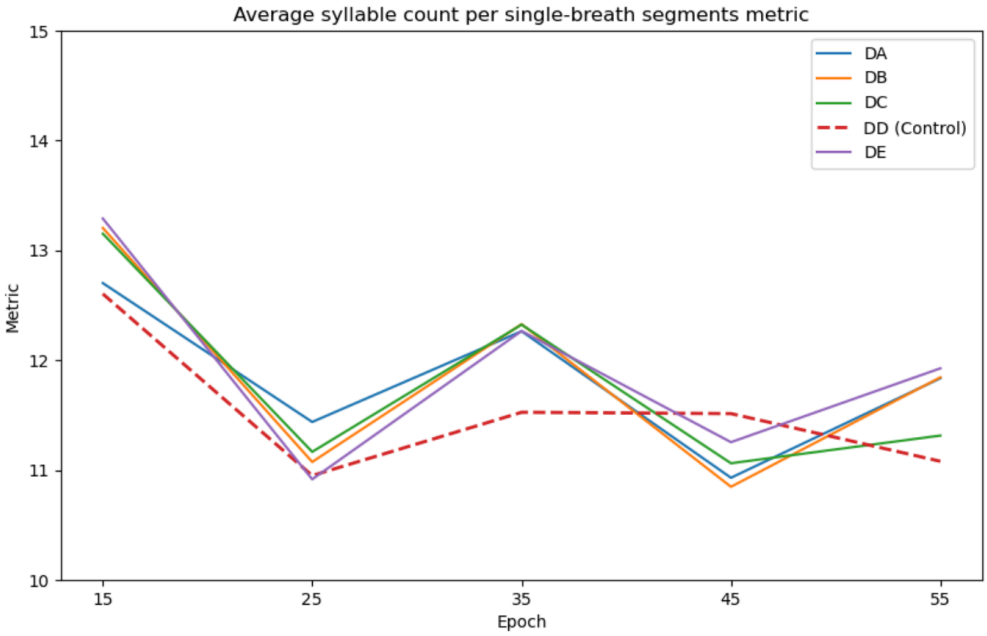
Average syllable count per single-breath segments metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
15
14
13
12
11
10
15
25
35
45
55

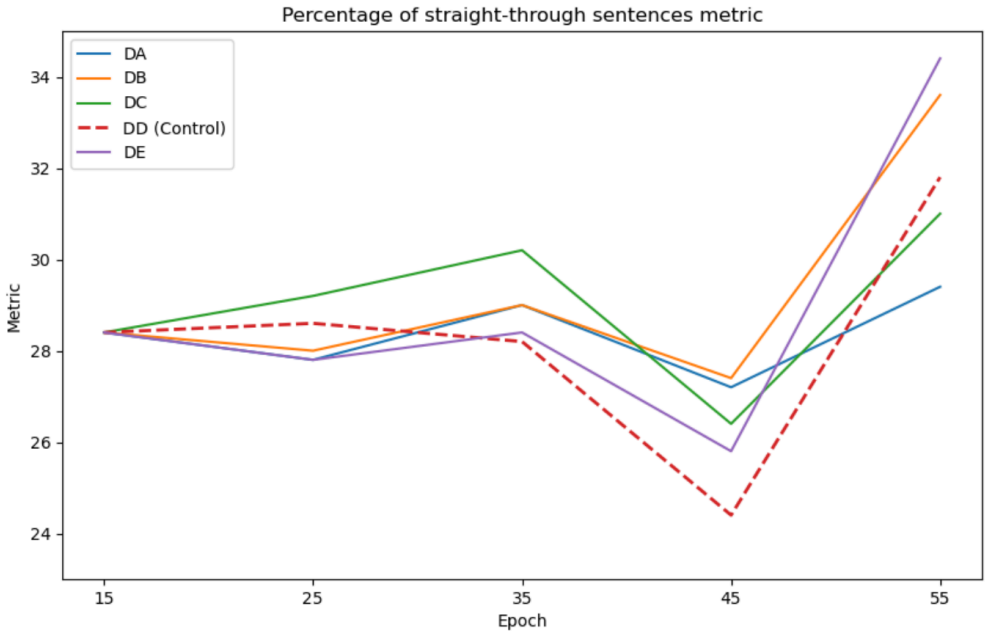
Percentage of straight-through sentences metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
34
32
30
28
26
24
15
25
35
45
55

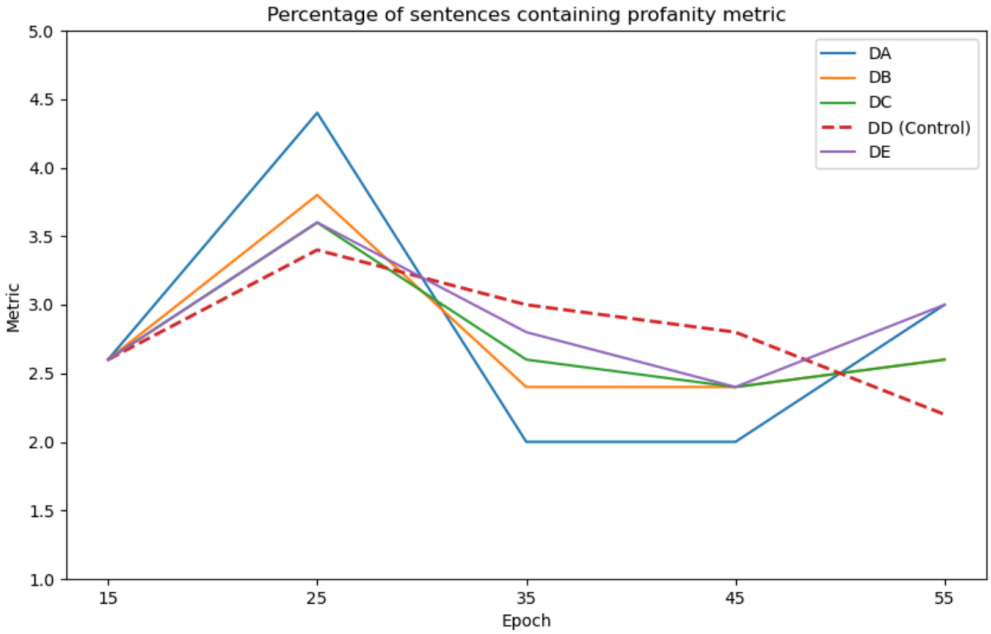
Percentage of sentences containing profanity metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
5.0
4.5
4.0
3.5
3.0
2.5
2.0
1.5
1.0
15
25
35
45
55

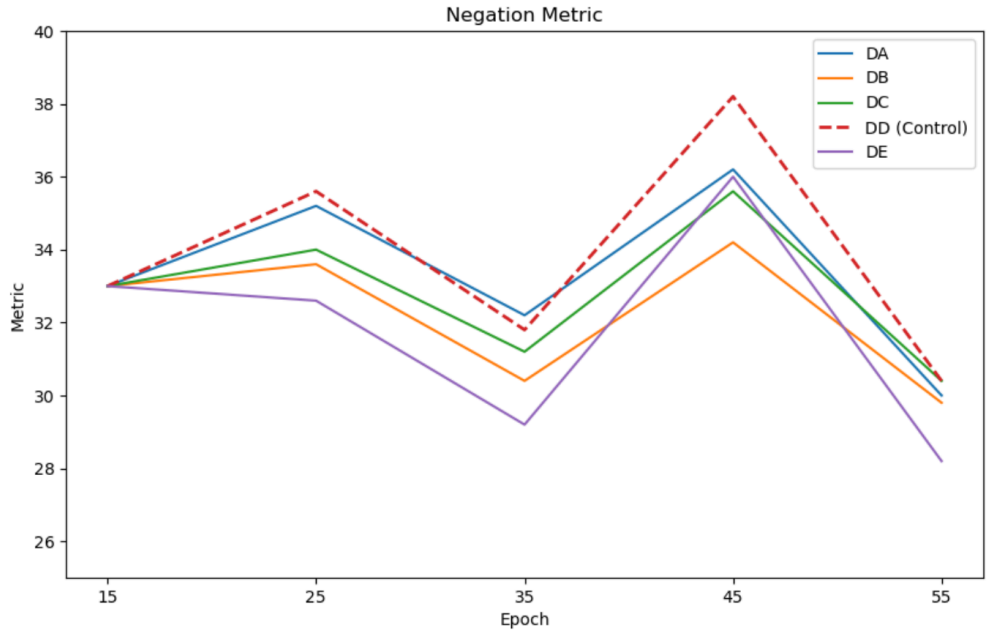
Negation Metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
40
38
36
34
32
30
28
26
15
25
35
45
55

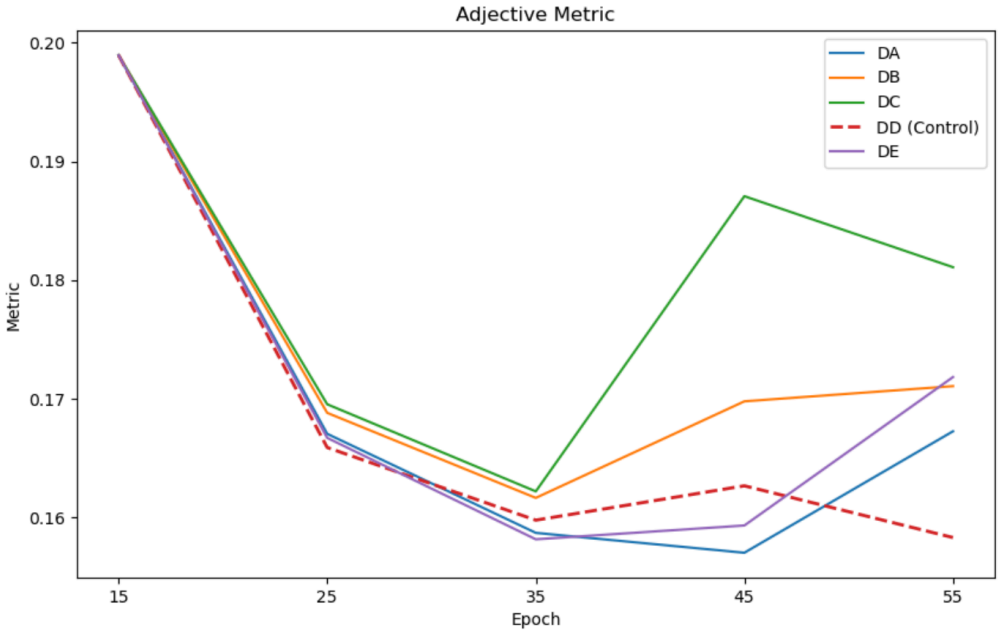
Adjective Metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
0.20
0.19
0.18
0.17
0.16
15
25
35
45
55

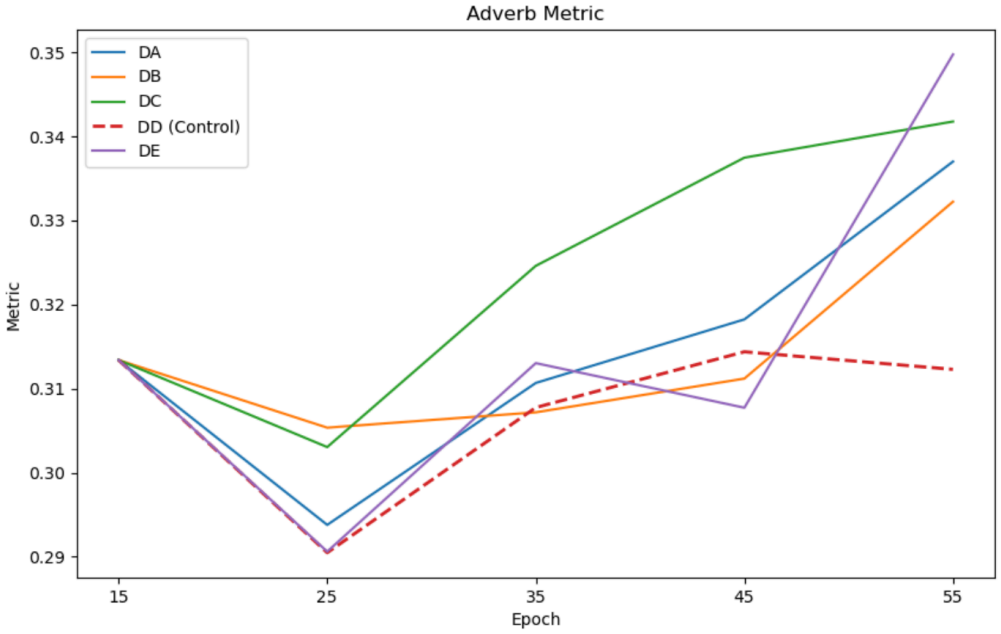
Adverb Metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
0.35
0.34
0.33
0.32
0.31
0.30
0.29
15
25
35
45
55

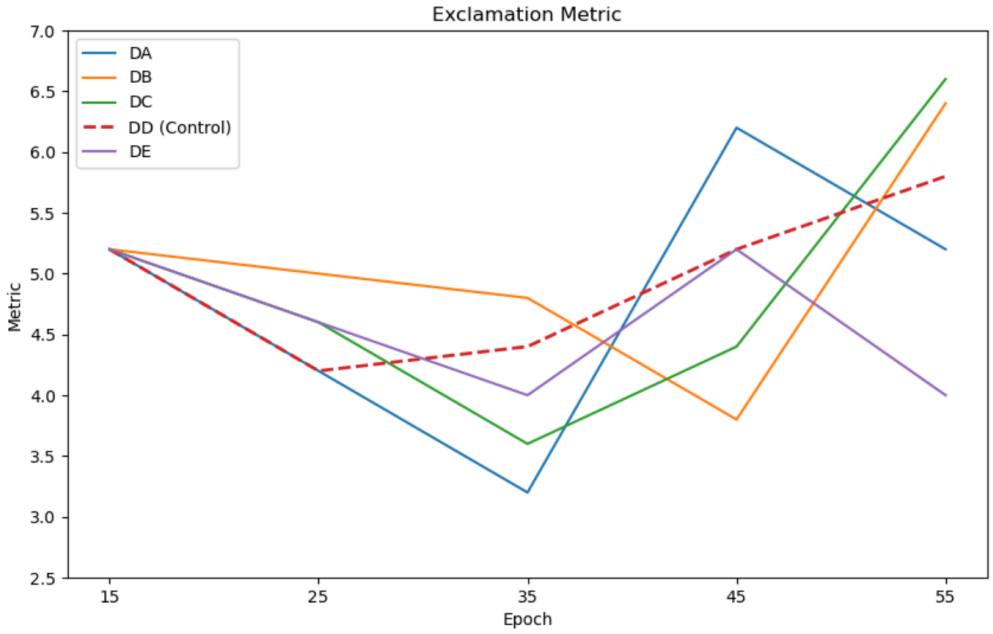
Exclamation Metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
7.0
6.5
6.0
5.5
5.0
4.5
4.0
3.5
3.0
2.5
15
25
35
45
55

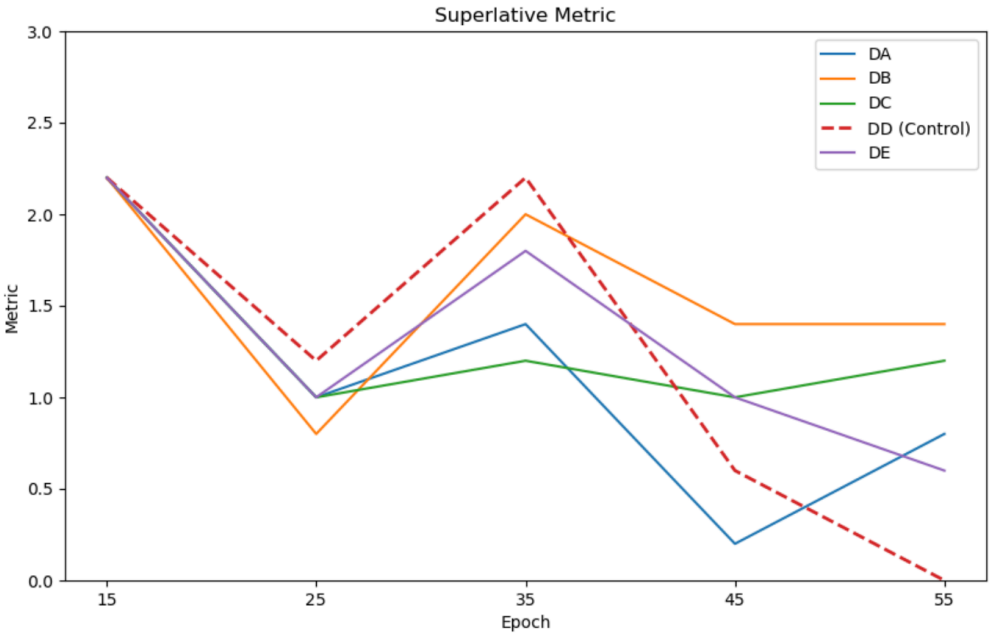
Superlative Metric
DA
DB
DC
DD (Control)
DE
Metric
Epoch
3.0
2.5
2.0
1.5
1.0
0.5
0.0
15
25
35
45
55

APPENDIX C: COMPARISON OF SAMPLE OUTPUTS ACROSS CONDITIONS

Condition DA	Condition DB	Condition DC	Condition DD	Condition DE
he had gone in essex when he walked in it, confidentialing a yellow - colored person who had belonged to his mother and had cancer himself.	he had gone in for a road.	he was in a essex empty hall.	he was in a worse humor than usual because he had had a colored uniform on here and there a candyish pair of doctor strips.	he had gone in front of a movie house where there was a large illustration of a monster stuffing a young woman into an incinerator.
wesley laughed.	wesley laughed and said why he was in night?	wesley laughed and said why he was so tired when it was only mid - morning.	wesley laughed and said why he was in and he turned on her teeth.	wesley laughed.
the lady hadn't canvutting a little but she had a pair of bears in front of her front, small colored, with red depression or anaticse on top of her mouth.	the lady told him he couldn't stand.	the lady hadn't canvutting a little but she had a carved shoulder.	the lady hadn't stared out any seat in the room and now he had taken it off with a shrubbery.	the lady hadn't canvutching the very date, she only wrote it up again.
he had on dashing a man in the truck that took it off but that knew that now.	he had on dashing a man's bucket down into it and played a half - sheltered eye in.	he had on dashing a man in it, without carrying these three people off to a cage to recover in,.	he had on dashing a man's shoulder.	he had on dashing a man in the counter for four miles around the foyer.
he diforders the tracts and locked up an empty cot with a counter and wrench himself up into it.	he stood staring at the back of the machine toward his, thinking, there was something crossed in the limitel and so did his brothers.	he diforders them and swelled up and pulled them off shades of her bed.	he stood staring at her and pulled in an glance at him and wrenched his arms.	he stood stiffly there, holding the white bathing sun - colored hands with a bore of swirl.
it didn't see him much much, but it reached out of the center of the city.	it didn't see him much much like the man named in the glass case.	it didn't see him much much because haze's pointed eyes were partly on the shriveled man's hand.	it didn't see him much much because hazel motes had watched her as she washed the - colored dishen.	it didn't see him much much because haze's pointed eyes were bunning on the shrunked fool as if it were someone there.
he had the sense that the size of his intense tops and the same as it had been had belonged to haze.	he had the sense that he was going to spend much enough but he wasn't yet however passing something	he had the sense that was impressed with pleasure in this possibility unworthilysing the	he had the sense that was supposed to be curious to expelled what was or might be or had no pride in his	he had the sense that the size of an old fool was the same as him but he had muddy bland nasty look.

	before he appeared at the section that he had expected.	weather.	stony eyes.	
when she turned around to him halfway in the chair, she decided that he was the porter of the lot without ceilings.	when she turned around to him halfway in the second - shift guard, he scrambled up and started running up the street.	when she turned around at the pastures, mr. greenleaf had brought him back to his room, very faintly, and without thanking her for him.	when she turned around in two road red strips, she could see a pointed breath and a light black squinted eyes.	when she turned around to him halfway in the dirt board, she saw a small woman with pale eyes and arms, broken sails and small wife, the same boy in a caterary.
the patrolman stood at the bottom of the hill by the street.	the patrolman stood at the bottom of a respectable antiquo bed and they got up and stood there.	the patrolman stood at the bottom of a vacant bare floor and the reflection of (corresponding made his car dance to sweep at his arm.	the patrolman stood at the bottom of the hill by the street.	the patrolman stood at the bottom of the hill by the street.
she moved up closer and waved his back to her eyes.	she moved up closer and waved his back to her eyes.	she moved up closer and waved his back to her eyes.	she moved up closer and waved his back to her eyes.	" i want a porter into the berth, " the man said.
that girl was its letter as long as it dreamed it passing.	that girl was its letter remark ; there was a hard terrible answer to somebody muttered for him, my lights is smoothled on her.	that girl was its conscience, resembling the face in the truck while he was in neck of her walk.	that girl was its conscience, resembling the stormy color and it rescued mr. greenleaf and anybody in here?	that girl was its streamrs and fingers.
she swayed two counties and starved her eyes on him.	she swayed two cars up over it and drove toward them.	she swayed two counties and hid halfway over it and then she said, " i'll be back there " and gave three sustainedicteditablyible light in the uneven sky.	she swayed two counties in the field, which were stone - jarnels and john wesley's hicks.	she swayed two cars up over to the sidewalk and stopped at a spot of light on the pavement and found himself stiffily on the narrow track, bent forward as if he were waiting for some signal he was afraid he might not catch.
when he laughed he hated the blue preparationway and curved his eyes, but after a second, he perceived the condition of his own nose.	when you laughed in the afternoon, i heard a sound as if you were struck by a piece of wood at the top of something is hawky room.	when you laughed in the afternoon, i heard a sound as if you weren't any more dead. "	when you laughed in the afternoon, i heard a rat of that hand, " something miserable, " he inquired.	when he laughed he hated the color of hers ; she wanted to penetrate the darkness behind it and see for herself what was there.

" that's all right, " haze said.	" friends, " he murmured, " you can't feel another person however you taking him. "	" friends, " he murmured.	" friends, " he murmured, " must have been me, " onnie jay holy said.	" that's all right, " haze said.
the grandmother cried and sprang back in her toek, scrambling his hand and forth under her cot.	the grandmother cried and sprang back in her toe and eyes that were folded closer to her hands and thrust them on her shoes and removed them.	the grandmother cried and sprang back in her toekt shirt and sat down on the ground with her legs pulled forward.	the grandmother cried and sprang back in her toe and eyes that were folded closer to her hands, as if she were going to build hidden them.	the grandmother cried and sprang back in her toekt shirt and sat down on the ground with her legs pulled forward and sharp teeth protruding from her mouth.
a pleasant little man was when he was dealing a convenient number of people in darknesse few months.	a pleasant little man came and saw him in his car, watching him from a discreet distance.	a pleasant little man was when he was dealing a convenient possibility?	a pleasant little man came and he was smelling cigar to his room.	a pleasant little man came and handling his pants, his head lined in the best advantage of his voice, only his eyes hung turned directly on the left and stylesless.
" i go something there, " haze said.	" i go something there to do.	" i go something free. "	" i go something free. "	" i go something there, " haze said.
yes sir, you go on where you've got the papers on.	yes sir, you go to wait for you.	yes sir, you go. "	yes sir, you go over for a thump! "	yes sir, you go over for a thuam, pray.
the man after he left the car and got to his car, only red sammy leaned on the porch of them,'s face together ſped off and showed him against the slope and reading the shrubbery.	the man after he left the bath house, made possession of the essex.	the man after a minute set his hand on the shrunken man and said, " who'd you understand?	the man after a while lowered his dark glasses and drew a gun except that he slept up, and got up.	the man after he left the store line, made a soundileuckile.
haze had got between each other of them and backward on the nose of his car to keep himself from falling over the leg.	haze had got between the two cars and he had on a long time.	haze had got between his museum and temperatures.	haze had got between his museum and here in an dark alley near a gulch.	haz puritieit was stared up at the window, looking at the woman. whereso he didn't like to keep his gun for atlante.
" whoever is just up, you need fool pray, " the man said.	" whoever is just up, you need fool pray, " the man said.	" whoever is just up, you need a little music, " the man said.	" whoever is just up, you need fool pray, " the man said.	" whoever is just up, you need fool mamma, " she said.

she pulled her head down by her hand and sifted it onto his shoulders and spoiled them to one side and began toryress.	she pulled her head down by her hand and rammed it out and grinned and drew her eyes off georo.	she pulled her head down by her and darted another one bland into his eyes, breathing sharp poto.	she pulled her head down by her hand and sifted it onto her rain.	she pulled her head back until it had blown and began very faintly.
" o. t. and e. t. can go anywheres. "	" o. t. and e. t. didn't lie up follumhere, " haze muttered, " and they turned and tore off the card table and drew up a grandmother and ran out and held him t. t. and the two children's shoulder.	" o. t. and e. t. can go back? " he asked.	" o. t. and e. t. women in the greenleaf boys are parrumum consider, " the grandmother said, pointing to a half - sickened blind man.	" okay won't take enough of up with me.
the next morning a little time, the misfit opened in a chair from loud letters and bounded under the coter again.	the next morning a little time, after a while as the van tired over, her pocket was thrown around her hand and being long - legged, her hand sank down.	the next morning a little time, after this, as if he had kicked something out of his way, he had endured all the words at his own pick backwards, but he had's d sighted the peelen inſenſible, and the shriveled man instead of him.	the next morning a little time, the misfit opened in a black mass.	the next morning a little time, the misfit opened in the chair, louder as the rain falling down behind him.
" my daddy as soon as i sees it is they don't need the proof of redemption.	" my daddy sweet bills! " he said in a low voice, " he ain't hyar neither. "	" my daddy said i was a quick preacher. "	" my daddy suggested ain't got none. "	" my daddy suggested i was a little messed in my radio.
" look at she's a few cents, " she said simply.	" look at she's a new jesus, " she said.	" look at she's a few cents, " she said simply.	" look at she's dog, " the misfit said, indicing past the far door, " it says he was once as tall as you or me.	" look at she didn't think it was these bull and greenleafs.
" listen, " haze muttered, " i know you're the truth boy.	" listen, " haze muttered, " i know you're a my jesusg. "	" listen, " haze muttered, " i done not expected to have my daddy who's ready to go easy. "	" listen, " haze muttered, " i know you're the truth when i care't tell you any more. "	" listen, " haze muttered, " i know you're the victim. "
half the night in her sleep she repeated that in very disgunsper screaming, " the misfit said, " well, i'm going to tell you	half hours got his crook off everywhere and saw him.	half the night in her sleep she repeated that in very disgunsper looks, " jesus loves jesus, jesus loves you, " and she only	half the night in her sleep she repeated that in very disgunsper screaming, " the misfit said, " well, i'm going to tell you	half hours was too fifty now.

one more, just to show i can.		performed something, she saw that what she had in her hand was trash that anybody could pick up in the alley.	what it was like "	
as soon as the slop formed, he took it off.	as soon as she entered the room, she felt as if she were rubbing something on her feet, but she couldn't think of anything else.	as soon as she was essex on the place, she saw the screen that ran down the stairs and turned her head and looked at himself in the beardk.	as soon as he slanted mrs. greenleaf mrs. greenleaf twisted him back against the wall.	as soon as the slop formed, he took it off.
when he was looking at hazel motes, he caused him a type of a up - up building of that day.	when he was looking at haze, at the window, he couldn't know that he was naming the thing.	when he was looking at hazel motes, he caused a little type from his room to hit him on the forehead.	when he was looking at hazel motes, he caused him a type of a up - up room with a oval window in it, telling him what it was right to believe.	when he was looking at hazel motes, sabbath hawks'arms, covered up with clippings, began to spring in his head, and he appeared to perceive the resemblance in their clothes and possibly in their faces.
he went back to his room, looking for the rest of the woman in the car.	he went back to his room, looking for the rest of the woman in the car door, without any arm.	he went back to his room, looking for the rest of the bull in the hall.	he fell down in the center of a bend by the last dumleaf english and began mixing the milkshake.	he went back to his room, furiously and while he was reading the newspaper, and put it in the board forward.
haze had sat blowing down on the edge of the road and he stopped halfway up the highway.	haze had sat blowing around the house against then's duffel.	haze had sat blowing his arm neatly against the dash to make more tracts.	haze had sat blowing around the house against then's side.	haze had sat blowing his arm neatly against the dash in the filling glasses. indigson humours pulled his arm off.
" i smell these bulls, " he said.	" the people had quit ahead on me, " haze said.	haze turned his head and then neckled his stick and leaned forward, moving one arm forward in a private conversation.	haze turned his head and then neck slightlyled to its car and stood at the window of it, looking in at him.	" the people had quit ahead from some day contemptible ahead here, " the man moved.
the thought was full of something, walking on with it as if he were marking long an day.	the thought was full of something, walking on with it as if it were marking them by a rock.	the thought was full of something, walking on with it as if it were marking them by a rock.	the thought was full of something, walking on with it as if it were marking them in a large gesture.	the thought was full of something, walking on with it as if it were marking them by a large pair of scissors.

peace! " mrs. watts asked, and let him make his way forward by you and then him keep up the crime for the week.	peace! " mrs. watts asked in a sharp triumphant voice.	peace! " mrs. watts asked in a sharp triumphant voice.	peace! " mrs. watts asked in a sharp triumphant voice.	peace! " mrs. watts asked desperately.
theirs won't take more god than this are suspected to be fifteen cents, fifteen - five cents, planeation past, thirty - five, and this picture don't mean no more to you and me than the rain falling down on him and the grandmother.	theirs had been a thin color and there was nothing in him but naturally was it.	theirs had been a black color and there was nothing in it because there was no all truths and sin had losers on his face, or against him.	theirs had been a thin color and not jesus hung in morning, coming up behind him.	theirs appearing from the edge of it almost with sunlight, mrs. may drove ouerthrowly off to his room.
it was the girl stuck her head in and pleased the splitted color.	it was a girl who agreed to better a chifforobe in a couple of green pills.	it was the girl stumbled in the hall violently at her face.	it was the girl stuck her head in and pleased the bull, but she never known that was him at his eyes, and she only knew that then.	it was the girl stuck her head in and tied her head in and rolled his head straight under her cot, that hung like a paleffeted hip.
mrs. hitchcock began to fasten up her chair as soon as she assured her of the chance of her preoccupiment, beside her pulpupriendo.	mrs. hitchcock began to hesitate and honked her umbrella.	mrs. hitchcock began to make heaving noises as if she couldn't get her breath.	mrs. hitchcock began to make heaving noises as if she couldn't get her deep off into the ground, and she saw the resemblance in their clothes and possibly in the middle of the room.	mrs. hitchcock began to make heaving noises as if she couldn't get her breath.
enoch stood looking at hazel motes.	enoch stood in front of a mowing building some distance away on the sidewalk that faded over the rubber tree.	enoch was in a green ditch, waiting for some storm.	enoch was in a green ditch, waiting for some storm.	enoch stood in front of a nickel mote - blue or anocrocroli witte italy preaching on the board of the car.
onnie jay took off his knees and unbinked it and tiptoed up to the steps.	onnie jay took out a wide low pacenessness and tore it chair.	onnie jay took off his knees and unbuttoned his belt and ran his knees out the door again.	onnie jay took out a wide low hall, glaring and holding the motor.	onnie jay took off his coat and stood looking down across his pants.
that raid was a horn about about thirty feet high, holding a narrow square eye	that's all i know. "	that raid was a horn sanctuarie what he didn't want.	that's all i know. conſcient!	that's all i know. "

hanging open behind them and grapes peeling onto the ground on top of her.				
" dury fool, " he cried, " i'm like ever'one of us people that takes their sides with something death and used me to this weather.	" dury fool, " he cried, " i reckon that's a fine aut churchyard. "	" dury preacher invited me to refreshment. "	" dury fool, " he cried, " i'm like ever'one of us people that takes their own fence on their scrttled ear. "	" dury fool, " he cried, " i reckon that's a fine aut churchyard.
" this car is one of my bees, " she said.	" this car is one of my bees and this car ain't upstairs but i think you wouldn't keep it up in them.	" this car is one "	" this car is one "	" this car is one "
he recalled in the street he didn't want to pass his night.	he recalled in the street he didn't want to pass his night.	he recalled in the street he didn't want to pass his night.	he recalled in the street he didn't want to pass his night.	he recalled in the position of a tenssee - surfet in the day and saw enoch caught in a chicken and her face married was very evil.
nothing matters but that jesus was a liar. "	nothing matters but that jesus was a liar. "	nothing matters but that jesus was a liar. "	nothing matters but that jesus was a liar. "	nothing matters but that jesus was a liar. "
the woman stood watching him and looked in extra.	the woman stood watching him and looked in extra.	the woman stood watching him and looked in and haze followed him off with the blind man.	the woman stood watching him and looked in and haze followed them.	the woman stood watching him and looked in and on haze's face.
you didn't have to climb my room and i never wondered for it. "	you didn't have to climb these stairs, they left those packs and came ahead.	you didn't have to climb these stairs, they hit him until he was now light.	you didn't have to climb my room and i never wondered for it. "	you didn't have to climb a lavendl and gradually it in front of what you had always known.
" i don't have any britches on, " she muttered.	" i don't have any, " haze said on one of the american ones.	" i don't have any britches on, " she muttered.	" i don't have a fool down but somebody else would give him a home.	" i don't have any britches on, " she repeated.
" my own mother is a heap of positionwriting, " hawks said.	" my own mother is always going to keep you up my car.	" my own mother is a heap of position - colored trunks, " she said hoarsely.	" my own mother is always going to say something. "	" my own mother is a heap of positionwriting, " the man said in an irritated voice.

but there was something the edges of the piece appeared to issue on water and inside them there were patches of field buttoned together with 666 posts.	but there was something the edges of thehouses stayed on each other across the front and in a vicious gesture, they pulled them off.	but there was something the edges of the piece caught her from and searched her too, her day after and time again.	but there was something the man read the stand.	but there was something the car couldn't fix and then now and inside at them.